ProCurve Networking

ProCurve Switch 5400zl, 3500yl and 6200yl Series Reviewer's Guide

ProCurve Networking
HP Innovation

Reviewer's Guide

invent

Table of Contents

Executive Summary 4

Introduction 4

Focus 4

ProCurve Switch 5406zl Intelligent Edge (J8697A) 4

ProCurve Switch 5412zl Intelligent Edge (J8698A) 4

ProCurve Switch 3500yl-24G-PWR Intelligent Edge (J8692A) 4

ProCurve Switch 3500yl-48G-PWR Intelligent Edge (J8693A) 4

ProCurve Switch 6200yl-24G-mGBIC (J8992A) 4

Audience 4

Scope 4

Product Positioning 4

Overview 4

ProCurve Switch 5400z and 3500yl Series 5

ProCurve Switch 6200yl 6

ProVision ASIC Architecture 6

Inside the ProVision ASIC Architecture 7

Classification and Lookup 7

Policy Enforcement Engine 7

Network Switch Engine Programmability 8

Fabric Interface 8

ProVision ASIC CPU 8

Fabric ASIC 8

Management Subsystem 9

Advanced Capabilities of the Product Family 9

Versatile Intelligent Port (VIP) 9

ProVision Hardware Resiliency 9

ProCurve Switch 5400zl Series 9

5400zl Chassis Layout 10

Power Supplies 12

Fan Tray 13

zl Modules 13

Management Module 13

Line Interface Modules 13

Power Supply Configurations 14

Specifications 16

Management Module 16

Line Interface Modules 17

ProCurve Switch 3500yl Series 20

Specifications 21

Additional Line Interface Module 22

ProCurve Switch 6200yl 22

Additional Line Interface Module 23

Overview of Features and Benefits 24

Feature Set Summary 26

Security Features 26

QoS Functions 27

Convergence 27

Layer 2 switching 27

Bridging Protocols 27

Routing Protocols 28
Multicast Protocols 28
Management 28
Standards and Protocols 29
Performance and Capacity 30
Capacity and Performance Features Comparison 30
Per-port Buffer Sizes 30
Routing and Forwarding Tables 31
Optimizing the 10-GbE Port Configuration 31
Throughput and Latency Performance Data 33
10 Gigabit Performance Traffic Patterns 34
Throughput Test 34
Latency Measurements 34
Comparison of 5400zl with 5300xl and 8100fl 35
Services & Support 36
Lifetime Software Updates (Best in the Industry) 36
Lifetime Warranty (Best in the Industry) 36
Free Telephone Support 36
Optional Support Services 36
Appendix A: Premium License 38
Intelligent Edge and Premium License 38
Appendix B: Policy Enforcement Engine 39
Policy Enforcement Engine Benefits 39
Wire-speed Performance for ACLs 40
Appendix C: Power over Ethernet 40
PoE device types 40
Power delivery options 41
PoE negotiation 41
Additional PoE Power – External Supplies 42
Support for Pre-802.3af Standard Powered Devices 42
Appendix D: PIM Sparse Mode 42
Appendix E: LLDP-MED 43
Appendix F: Virus Throttle Security 44
Response options 45
Sensitivity 46
Connection-rate ACL 46
Appendix G: VRRP 46
XRRP support on 5300xl switch 47
Appendix H: OSPF Equal Cost Multipath 47
Appendix I: Troubleshooting 48
LED Status Indicators for 5400zl 48
LED Status Indicators for 3500yl and 6200yl 51
Part Numbers and Field Replaceable Units 52

Executive Summary

ProCurve Networking has an extensive line of networking products built around the concept of the ProCurve Adaptive Edge Architecture (AEA) that provides the security, mobility, and convergence capabilities that businesses demand while giving IT administrators the ability to adapt to the changing needs of their organizations and controlling their infrastructure centrally.

This guide describes the ProCurve Switch 5400zl, 3500yl and 6200yl series which build upon the principles of the Adaptive EDGE, bringing intelligence and performance to the Edge in a cost-effective way and providing a platform for delivering the full Intelligent Edge capability. The foundation for these switches is a purpose-built ProVision™ ASIC that allows the most demanding networking features, such as QoS and security, to be implemented in a scalable yet granular fashion. With a high performance architecture, 10-GbE capability, and programmable ASIC, these switches offer excellent investment protection, flexibility, and scalability.

Introduction

Focus

The current revision of this guide covers the following ProCurve switch products:

ProCurve Switch 5406zl Intelligent Edge (J8697A)

The ProCurve Switch 5406zl is a 4U switch with six zl line interface module slots designed to support up to 144 Gigabit or 24 10-Gigabit ports. The throughput of the 5406zl switch fabric is capable of up to 214 million (64 byte).packets per second (Mpps) with a backplane speed of 346 Gbps.

ProCurve Switch 5412zl Intelligent Edge (J8698A)

The ProCurve Switch 5412zl is a 7U switch with twelve zl line interface module slots designed to support up to 288 Gigabit or 48 10-Gigabit ports. The throughput of the 5412zl switch fabric is capable of up to 428 Mpps with a backplane speed of 692 Gbps.

ProCurve Switch 3500yl-24G-PWR Intelligent Edge (J8692A)

The ProCurve Switch 3500yl-24G-PWR is a 1U switch designed to support 20 10/100/1000 PoE ports, four dual personality ports, and four 10-Gigabit ports. The throughput of the 3500yl-24G-PWR switch fabric is capable of up to 74 Mpps with a backplane speed of 115 Gbps.

ProCurve Switch 3500yl-48G-PWR Intelligent Edge (J8693A)

The ProCurve Switch 3500yl-24G-PWR is a 1U switch designed to support 44 10/100/1000 PoE ports, four dual personality ports, and four 10-Gigabit ports. The throughput of the 3500yl-24G-PWR switch fabric is capable of up to 110 Mpps with a backplane speed of 173 Gbps.

ProCurve Switch 6200yl-24G-mGBIC (J8992A)

The ProCurve Switch 6200yl-24G-mGBIC is a 1U switch designed to support up to 24 mini-GBIC ports and up to four 10-Gigabit ports. The 6200yl-24G-mGBIC is designed to be deployed as an aggregator of traffic from the edge to the core of the network. The throughput of the 6200yl-24G-mGBIC switch fabric is capable of up to 74 Mpps with a backplane speed of 115 Gbps.

Audience

This guide is primarily written for technical evaluators and product reviewers of networking equipment and solutions.

Scope

This guide provides detailed information and specifications about the ProCurve Switch 5400zl, 3500yl and 6200yl series products, with the assumption that details about standard protocols can be referenced externally by those familiar to general networking. Those technologies that are relatively new will be covered in more detail than the older, less broadly familiar, technologies.

Product Positioning

Overview

A widely used method for segmenting the areas in which switches are installed calls for three different classifications: access, distribution and core. Access switches provide aggregation of

end nodes for connection to a distribution or core switch and are usually found in wiring closets. Distribution switches aggregate the links from access switches and possibly server farms. Distribution switches anchor the network in a building, or, for small networks, across a campus. Core switches provide the focal point of the local network, aggregating the distribution switches, providing connectivity for central site data centers, and providing connectivity in many cases to the external network.

Regardless of your network's architecture — whether you have adopted ProCurve's Adaptive EDGE architecture with its intelligence at the edge, or whether you use a traditional architecture that concentrates intelligence at the center of the network — you have a choice of wiring topologies to implement that architecture.

The intelligence and throughput of the ProCurve Switch 5400zl and 3500yl series make them suitable for applications at or near the edge of networks of all sizes, or in the center of networks that are small-to-moderate in size. The modular structure of the ProCurve Switch 5400zl series gives you the flexibility to tailor the switch to match your topology's requirements. The ProCurve Switch 6200yl, an advanced Layer 3 stackable, is designed to be deployed as an aggregator of traffic from the edge to the core of the network.

The foundation for all of these switches is a purpose-built, programmable ProVision ASIC that allows the most demanding networking features, such as QoS and security, to be implemented in a scalable yet granular fashion. A high capacity switch fabric, based on the most recent ProVision ASIC architecture, is integrated with the switch backplane. The ProCurve Switch 5400zl, 3500yl and 6200yl series have been designed as a product family, using the ProVision ASICs and software, providing consistency and scalability across the family.

The ProVision ASICs are aimed at accomplishing several objectives:

- Provide a great engineering balance between feature capabilities, performance, and price
- Bring sophisticated control features to the edge of the network
- Have programmable features that allow future requirements to be accommodated in the ASICs

Key features of the ProCurve Switch 5400zl, 3500yl and 6200yl series include:

- Performance—High capacity switch fabric (from 115 Gbps to 692 Gbps backplane speed), bandwidth shaping and control, Quality of Service, L2 and L3 jumbo frames
- Security—ACLs (per-port or identity-driven), virus throttle, switch CPU protection, detection of malicious attacks, DHCP protection, BPDU port protection, Dynamic ARP protection, Dynamic IP lockdown (target availability 2Q 2007), IP and MAC lockdown/lockout, 802.1X, Web, and MAC user authentication, management access control (SSH, SSL, TACACS+, secure FTP)
- Operational flexibility—High-port density in 4U and 7U form factors, Versatile Intelligent Ports (10/100/1000, PoE-enabled), power supply choices for optimum PoE
- Resiliency—Redundant power supplies, hot-swappable/hot-insertable modules, MSTP, switch meshing, VRRP, OSPF-ECMP
- IP Routing—RIPv1, RIPv2, OSPF, PIM-SM, PIM-DM, static routes
- Diagnostic—Remote intelligent mirroring, loopback interface, UDLD, sFlow support
- Investment protection—Upgradeable management engine and CPU, upgradeable to Premium License feature group, add-in modules and power supplies, Versatile Intelligent Ports, programmable ASICs

ProCurve Switch 5400z and 3500yl Series

The ProCurve Switch 5400zl and 3500yl series consists of the most advanced intelligent edge switches in the ProCurve Networking product line. The 5400zl series includes 6-slot and 12-slot chassis and associated zl modules and bundles, and the 3500yl series includes 24-port and 48-port stackables.

With a variety of Gigabit interfaces, integrated Power over Ethernet (PoE) on all 10/100/1000Base-T ports, 10-GbE capability, and a choice of form factors, the 5400zl and 3500yl switches offer excellent investment protection, flexibility, and scalability, as well as ease of deployment, operation, and maintenance.

The 5400zl and 3500yl switches are targeted as enterprise class wiring closet switches—designed for low cost with a choice of medium to high port density. Voice, video and data ready,

the ProCurve Switch 5400zl and 3500yl series offer extensive prioritization features that bring full convergence down to the desktop. Integrated PoE minimizes wiring requirements for VoIP phones and wireless Access Points. The ProCurve Switch 5400zl and 3500yl series provide the fine-grained security at the edge of the network to lockout external threats, yet provides appropriate access to employees and guests. Collectively, these features make the 5400zl and 3500yl switches well-suited for the access tier.

For some customers, the Layer 3 features and redundant power supply features of the ProCurve Switch 5400zl series also make it well-suited as a distribution switch. New to the 5400zl and 3500yl series is the optional deployment of a *Premium License* feature group as opposed to the standard offerings of the *Intelligent Edge* feature group. Premium License contains all of the Intelligent Edge features plus additional features oriented towards routing that are useful when deploying the switch at the distribution level of the network. The Premium License feature group includes protocols such as PIM-SM, PIM-DM, OSPF-ECMP and VRRP.

ProCurve Switch 6200yl

The ProCurve Switch 6200yl-24G-mGBIC is an advanced Layer 3 stackable in 1U height. It has 24 mini-GBIC slots and an expansion slot for an optional 4-port 10-GbE module. Designed to be deployed as an aggregator of traffic from the edge to the core of the network, this switch supports a variety of Gigabit mini-GBICs, such as SX, LX, LH, and 1000Base-T. The *Premium License* feature group is standard on the ProCurve Switch 6200yl-24G-mGBIC.

Like the 5400zl and 3500yl series switches, the foundation for this switch is a purpose-built ProVision ASIC that allows the most demanding networking features, such as QoS and security, to be implemented in a scalable yet granular fashion. With its high performance architecture, 10-GbE capability, and programmable ASIC, this switch offers excellent investment protection, flexibility, and scalability.

ProVision ASIC Architecture

The ProVision application-specific integrated circuit (ASIC) architecture is the latest generation ProCurve ASIC technology and is used in the ProCurve Switch 5400zl, 3500yl and 6200yl series. The ProVision ASIC architecture consists of multiple network chips interconnected by an active crossbar consisting of the fabric chip. A network chip is implemented on each of the various line interface modules (also known as a line cards). The fabric chip is integrated on the switch backplane which helps reduce the overall size of the switch.

Each network chip represents a node in the system with “links” connecting to the interconnect fabric. Each link provides approximately 28.8 Gbps of data bandwidth. In addition, a management module with a dedicated CPU provides communications control between the network chips and fabric chips.

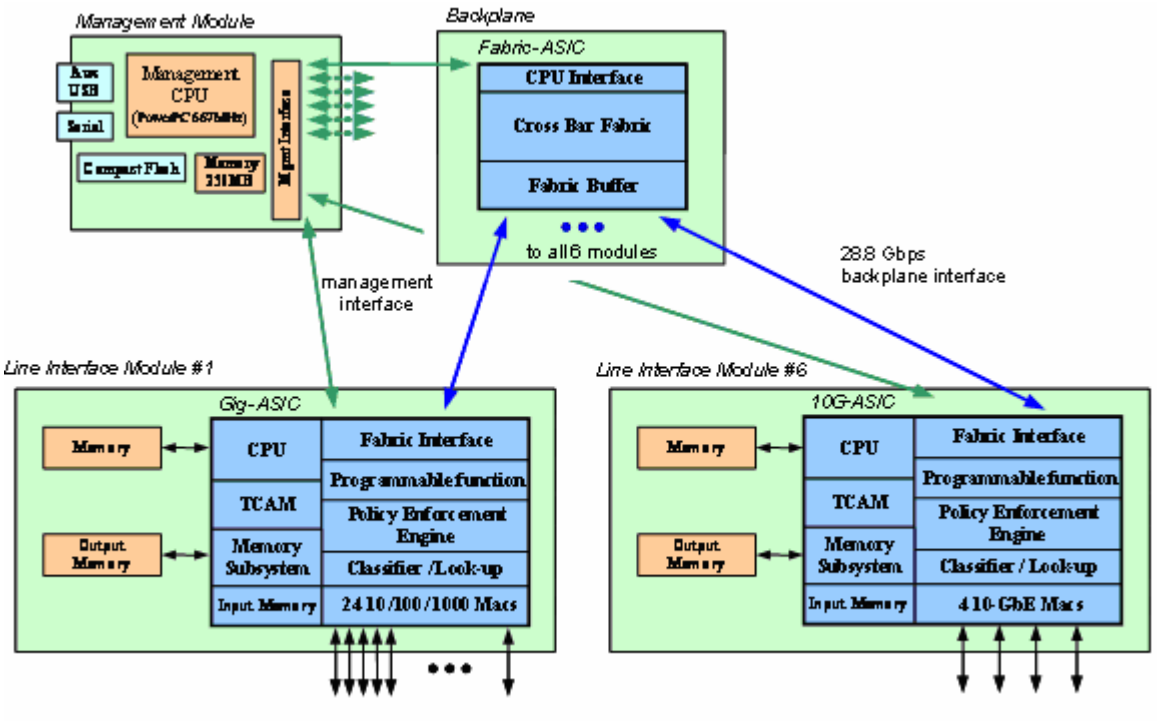

Figure 1. ProVision ASIC Architecture for 5406zl

The diagram above illustrates an example of the logical interconnection of the ProVision ASICs on the 6-slot 5406zl series switch. All of the key elements are connected to the active backplane. The active backplane contains the switch fabric and distributes power to all modules. The ProCurve Switch 5412zl, 3548yl, 3524yl and 6224yl have similar architectural components. The primary difference in the illustration would be the number of ports supported.

Inside the ProVision ASIC Architecture

Each line interface module contains a full ASIC-based Layer 3 routing switch engine as well as Layer 4 filtering and metering. These new ProVision ASICs are ProCurve's 4th generation switching ASICs. This network switch engine, in the ProVision ASICs, provides all the packet processing: Layer 2 and Layer 3 lookups, filtering and forwarding decisions, VLAN, trunking and priority queuing determinations. The ProVision ASIC on each line card contains its own CPU. These features of the ProVision ASIC are common for all products in the ProCurve Switch 5400zl, 3500yl and 6200yl series.

Classification and Lookup

When a packet first comes in, the classifier section determines the packet characteristics, its addresses, VLAN affiliation, any priority specification, etc. The packet is stored in input memory, lookups into the table memory are done to determine routing information, and a ProVision ASIC specific packet header is created for the packet with this information. This header is then forwarded to the Policy Enforcement Engine.

Policy Enforcement Engine

The ProVision ASICs on each line interface module contain the Policy Enforcement Engine. This engine provides fast packet classification to be applied to ACLs, QoS, Rate Limiting and some other features through an onboard TCAM. Some of the variables that can be used include source and destination IP addresses (can follow specific users), TCP/UDP port numbers and ranges (apply ACLs to an application that uses fixed port numbers or ranges). Over 14 different variables can be used to specify the packets to which ACL and QoS rules, rate limiting counters, and others are to be applied.

Partially implemented in the initial software release, the Policy Enforcement Engine will provide a common front end for the user interface to ACLs, QoS, rate limiting and some other services. In subsequent software releases for the switches, more features can take advantage of the

Policy Enforcement Engine to provide a powerful, flexible method for controlling the network environment. For example, traffic from a specific application can be raised in priority for some users, blocked for some other user, and limited in bandwidth for yet other users. After the policy enforcement engine, the header is then forwarded to the programmable section of the network switch engine.

Network Switch Engine Programmability

Each ProVision ASIC switch engine contains multiple programmable units making them truly Network Processor Units (NPU). One of the functions of the NPU is to analyze the header of each packet as it comes into the switch. The packet's addresses can be read with the switch making forwarding decisions based on this analysis. For example, if a packet's 802.1Q tag needs to be changed to re-map the packet priority, the ProVision ASIC needs to look at each packet to see if any particular one needs to be changed. This packet-by-packet processing has to occur very quickly to maintain overall wire-speed performance – a capability of the ProVision ASICs.

To broaden the flexibility of the ProVision ASICs, a programmable function is included for its packet processing. This NPU function, allows the ProCurve designers the opportunity to make some future changes or additions in the packet processing features of the ASIC by downloading new software to it. Thus, new features needing high performance ASIC processing can be accommodated, extending the useful life of the switch without the need to upgrade or replace the hardware. In the first release of the ProCurve Switch 5400zl, 3500yl and 6200yl series, the NPU function within the ProVision ASICs is totally unused, awaiting future upgrades.

The concept of adding the programmable functionality of the NPU within a switching ASIC was originally designed and implemented in the popular ProCurve Switch 4000M switch family introduced in 1998. The ProCurve's 5300xl programmable capability was a second generation design based on the original ProCurve Switch 4000M implementation. The programmable capability was used to give both the ProCurve Switch 4000M and Switch 5300xl new ASIC-related features well after initial release of those products. Customers with existing units could benefit from the new features through free software downloads. The customer's investment in the ProCurve Switch 4000M and 5300xl are preserved by providing new functionality not otherwise possible without the ASIC NPU programmability.

Being based on the ProCurve Switch 4000M and 5300xl implementations, the NPU capabilities of the ProVision ASICs used in the ProCurve Switch 5400zl, 3500yl and 6200yl series are a third generation design.

Fabric Interface

After the packet header leaves the programmable section, the header is forwarded to the Fabric Interface. The Fabric Interface makes final adjustments to the header based on priority information, multicast grouping, etc., and then uses this header to modify the actual packet header as necessary.

The Fabric Interface then negotiates with the destination ProVision ASICs for outbound packet buffer space. If congestion on the outbound port is present, WRED (weighted random early detection) can also be applied at this point as a congestion avoidance mechanism. Finally the ProVision ASICs Fabric Interface forwards the entire packet through the Fabric-ASIC to an awaiting output buffer on the ProVision ASICs that control the outbound port for the packet. Packet transfer from the ProVision ASICs to the Fabric-ASIC is accomplished using the 28.8 Gbps full duplex backplane connection, also managed by the Fabric Interface.

ProVision ASIC CPU

Each ProVision ASIC contain its own CPU for learning of Layer 2 nodes, packet sampling for the XRMON function, handling local MIB counters, and running other module related operations. Overall, the local CPU offloads the master CPU by providing a distributed approach to general housekeeping tasks associated with every packet. MIB variables, which need to be updated with each packet, can be done locally. The Layer 2 forwarding table is kept fresh through the use of this CPU. Other per-port protocols, such as Spanning Tree and LACP, are also run on this CPU. The local CPU, being a full-function microprocessor, allows functionality updates through future software releases.

Fabric ASIC

The Fabric ASIC, which is located on the backplane of the switch, provides the crossbar fabric for interconnecting the modules together. The use of a crossbar allows wire speed connections simultaneously from any module to any other module. As mentioned in the ProVision ASICs section, the connection between the Fabric-ASIC and each line interface module's ProVision ASIC is through a 28.8 Gbps full duplex link.

Management Subsystem

The management subsystem is responsible for overall switch management. The management subsystem consists of a CPU, flash memory to hold program code, processor memory for code execution, status LEDs and pushbuttons, a console interface, and other system support circuitry to interface and control each line interface module. In the case of the 5400zl series, the management subsystem is on a module that is removable/upgradeable. Each 5400zl series chassis requires one management module to function. For the ProCurve Switch 3500yl and 6200yl series, the management subsystem is an integrated component.

Advanced Capabilities of the Product Family

The ProCurve Switch 5400zl, 3500yl and 6200yl series product family includes a number of advanced capabilities that offer a highly reliable, robust chassis data environment that leads to increased network uptime, keeping overall network costs down.

Versatile Intelligent Port (VIP)

All Gigabit copper ports are PoE enabled and tie into the overall capabilities of the switch, such as the Policy Enforcement Engine. Having the VIP, any port can be connected to any place in the network. If a PC needs Gigabit connectivity it is there. If an AP needs PoE it is there. And if for some reason, these ports need to be swapped on the switch, that can be done by just moving the cables at the switch or patch panel.

The Versatile Intelligent Port provides a high degree of flexibility, makes troubleshooting easier, and provides a high degree of future proofing.

ProVision Hardware Resiliency

Many functions required in a switch have been implemented in the single ASIC on the module. What takes a number of chips in other vendor products is achieved in a single ProVision ASIC. This keeps the part count down raising overall reliability of the module a significant degree.

Another engineering aspect in the ASIC is hardware error detection with correction in software for the memory used by the switch. This includes the memory used for forwarding the network traffic such as the routing and forwarding tables, the Policy Enforcement Engine information, multicast tables, and other data structures. Traffic sent across the backplane uses a protocol to check that there is space available at the destination module so that fabric data is not lost.

ProCurve Switch 5400zl Series

The ProCurve Switch 5400zl series is a high-end Edge switch in the ProCurve product line. This switch series has been designed to fill the needs of a feature-oriented high-performance wiring closet switch. It can also be used as a low-to-medium distribution switch when used with *Premium License*, described later in this document.

The 5400zl series is the latest generation of managed chassis products from ProCurve. The 5400zl series follows in the evolution of ProCurve managed chassis switches, featuring Power over Ethernet (PoE) on every copper port, chassis flexibility at the stackable price range, and a scalable architecture.

The 5406zl and 5412zl each require a single Management module which initializes, controls, and monitors the various line modules and the switch fabric. The 5406zl has six slots available for population with line interface modules whereas the 5412zl has twelve slots available. The 5406zl and 5412zl support the same line interface modules. These line interface modules are the 24-port 10/100/1000, 20-port 10/100/1000 + 4-port mini-GBIC, 24-port mini-GBIC, 4-port 10-GbE X2, and 4-port 10-GbE CX4. The 5406zl supports up to two internal power supplies and the 5412zl supports up to four internal power supplies. The power supplies can be used for purposes of power redundancy or for additional PoE requirements.

The base configuration for the 5406zl (J8697A) includes a management module, Intelligent Edge software, and 6 open slots. Line interface modules can be added as needed for connectivity. The base configuration for the 5412zl (J8698A) includes a management module, Intelligent Edge software, and 12 open slots. Like the 5406zl, line interface modules can be added as needed for connectivity. Note that NO power supplies ship with base chassis products—they must be ordered separately.

The management module monitors the ambient temperature of the system. In the event the temperature exceeds a product specified threshold, an SNMP trap and event log entry is generated.

ProCurve Switch 5400zl Chassis

ProCurve Switch5406zl

(J8697A)

ProCurve Switch 5412zl

(J8698A)

ProCurve Switch 5400zl Bundles

ProCurve Switch 5406zl-48G

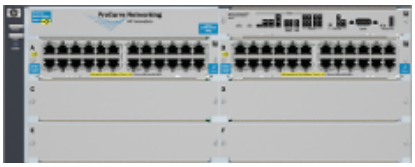

(J8699A)

ProCurve Switch 5412zl-96G

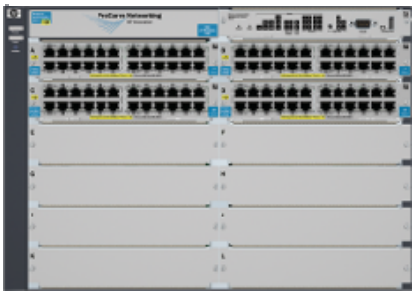

(J8700A)

Figure 2. ProCurve Switch 5400zl Chassis and Bundles

The *5406zl-48G Intelligent Edge (J8699A)* and *5412zl-96G Intelligent Edge (J8700A)* are preconfigured bundles that offer a pre-tested environment and a lower cost starter switch to which additional modules can be added. The major components of these preconfigured bundles are the following:

5406zl-48G Intelligent Edge	5412zl-96G Intelligent Edge
6-slot chassis	12-slot chassis
1 management module	1 management module
1 fabric module	2 fabric modules
48 10/100/1000 PoE ports	96 10/100/1000 PoE ports
4 open slots	8 open slots
1 fan tray (2 fans)	1 fan tray (4 fans)
1 internal 875W power supply	2 internal 875W power supplies
1 open power-supply slot	2 open power-supply slots
Intelligent Edge software	Intelligent Edge software

5400zl Chassis Layout

The 5406zl chassis is rack-mountable with a 4U height and the 5412zl is rack-mountable with a 7U height. The line interface modules are inserted in the front slots, labeled A through F on the 5406zl, and A through L on the 5412zl.

The Management module is removable/replaceable and occupies a dedicated slot in the front. The switch does not have to be powered-off to remove the management module.

However, when the Management module is removed, all ports lose communication and the system will be powered down.

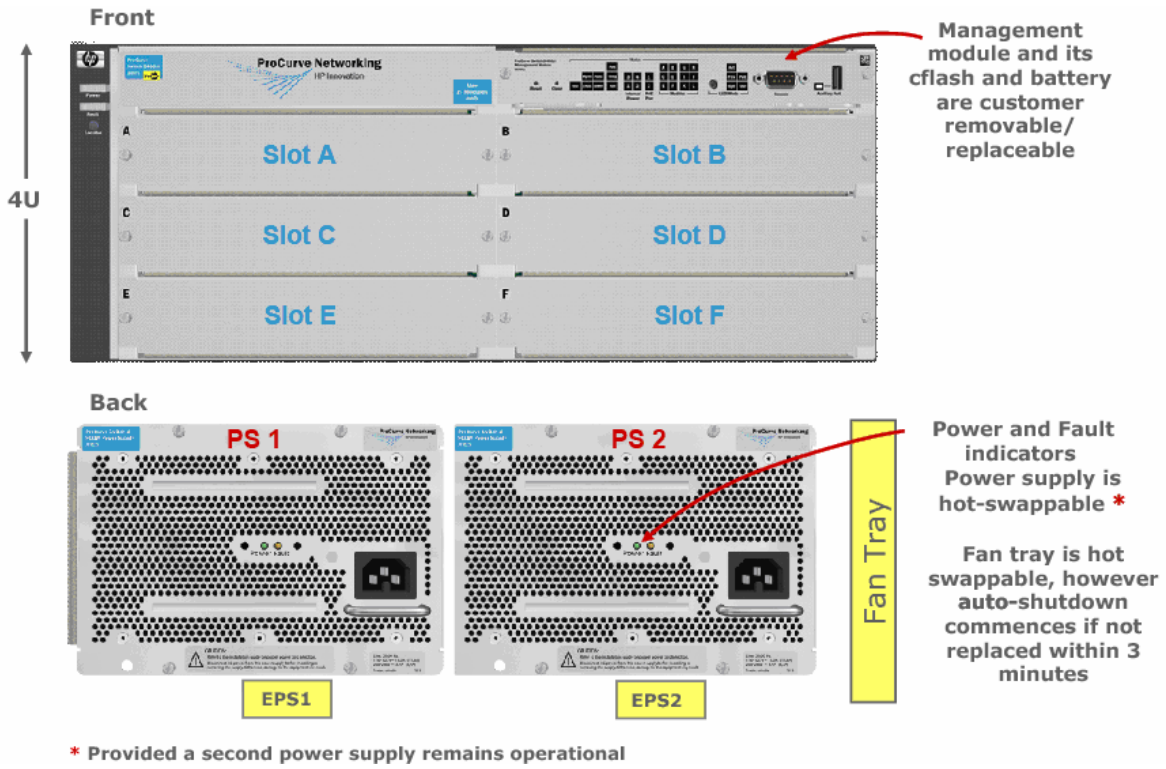

Figure 3. ProCurve Switch 5406zl Chassis Layout

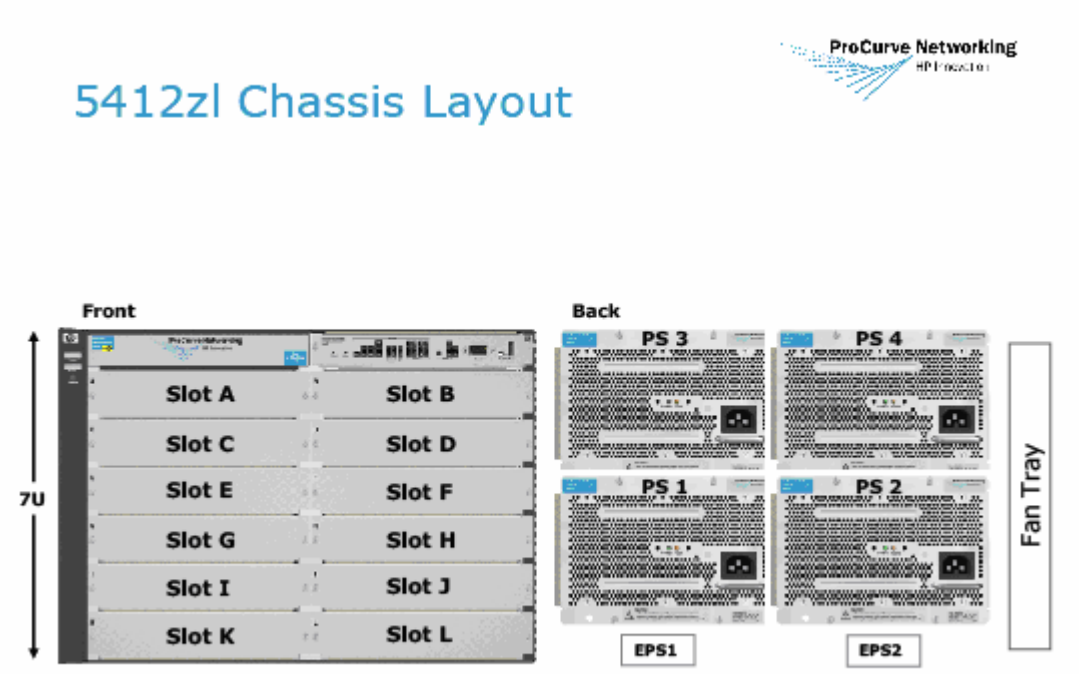

Figure 4. ProCurve Switch 5412zl Chassis Layout

The internal power supplies are inserted in the back slots. These slots are labeled PS1 and PS2 on the 5406zl and PS1 through PS4 on the 5412zl. A power supply is hot-swappable provided at

least one other power supply is operational. If the 5412zl has only two power supplies and one of them fails, then only the upper six slots (slots A through F) will receive power.

Power Supplies

There are two different power supplies available for the ProCurve Switch 5400zl series. Both power supplies provide system power (the power needed to run the switch itself) and PoE power (the power sent down the Ethernet cable to power the device at the other end). The only difference between the two supplies is the amount of PoE power available from the supply.

All ProCurve Switch 5400zl series chassis use the same power supplies. There are two internal power supplies available:

- 875W (110V/220V)
- 1500W (220V)

As indicated in the following figure, the internal power supply provides both system power as well as PoE. The ProCurve Switch 5400zl series, as well as the 3500yl series, can power any device that adheres to the IEEE 802.1af standard. In addition, devices using pre-standard Cisco PoE power can also be powered. The switches will automatically detect what type of power and how much is needed when a compatible device is plugged into the port.

It is recommended that you use the same power supply model for all power supplies installed in a given 5400zl series chassis.

Power Supply Types	System Power	PoE Power
875W Power Supply J8712A (110/220 VAC)	600W	273W
1500W Power Supply J8713A (220 VAC only)	600W	900W
ProCurve Switch zl Power Supply Shelf (J8714A)	0	Up to 1800W

875W Internal Power Supply

Figure 5. Power Supply Choices

The internal power supplies provide system power for all internal components (+12V output within +/- 5% tolerance, regulated) and power for PoE ports (-48V output within +/- 5% tolerance, regulated), which meets isolation and noise requirements of the IEEE 802.3af specification.

Both internal supplies have over-current, over-temperature and over-voltage protection, and integrated fans. Hot swapping is allowed, taking into account that disconnecting the power supplies may interrupt PoE operation.

An external power shelf, the ProCurve Switch zl Power Supply Shelf, is available to house up to two switch zl power supplies. The power shelf supplies can be connected to one or two switch zl's to provide additional PoE power, either to power additional PoE powered devices, and or provide more power for PoE power redundancy. The supplies in the power shelf are connected to the switch(es) via one or two 2m EPS cables. The PoE power provided is added to the internal power supply PoE power to figure the total amount of PoE power to the ports in the switch.

While the connectors on the EPS connector cable will fit the connectors on the 3500 switches, no power will flow from the power shelf. Extra PoE power for the 3500 switch can be obtained with the ProCurve 620 Redundant/External Power Supply.

Fan Tray

The fan tray assembly contains the cooling fans for the interior of the 5400zl chassis, but excludes the chassis power supplies, which have their own internal cooling fans. The 5406zl fan tray consists of 2 variable speed fans and the 5412zl fan tray consists of 4 variable speed fans. The fan speed is based on the sensed ambient temperature of the chassis.

The fan tray is mounted on the left side of the chassis (from a front-view perspective) and the fans draw air through ventilation holes to the left and blow the air out through ventilation holes to the right (side-to-side airflow).

The fan tray is hot swappable in the 5406zl and 5412zl. It can be removed and replaced without removing power from the switch. However, the new fan tray should be installed immediately after removing the old fan tray to avoid overheating and automatic switch shutdown within three minutes.

In the event of an individual fan failure, an SNMP trap and event log entry is generated.

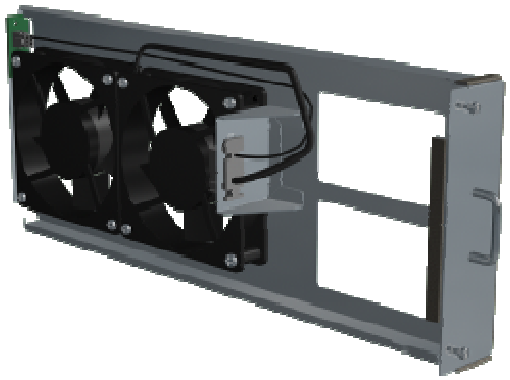

Figure 6. Fan Tray for 5406zl (J8697-60005)

zl Modules

Management Module

All configurations of the ProCurve Switch 5400zl series include a single Management module that oversees the operation of the line interface modules and switch fabric. The Management module incorporates an RS-232 serial port for local management and configuration. To connect to the console, a standard null modem cable is used.

Figure 7. ProCurve Switch 5400zl Series Management Module

Line Interface Modules

The ProCurve Switch 5400zl series supports a variety of popular line interface modules providing customers with the ability to change or scale their LAN links and adapt as the needs of the business environment change over time.

A variety of line interface modules are available for ProCurve Switch 5400zl series:

- ProCurve Switch 5400zl 24-port 10/100/1000 PoE Module (J8702A)
- ProCurve Switch 5400zl 20-port 10/100/1000 + 4-port Mini-GBIC Module (J8705A)
- ProCurve Switch 5400zl 24-port Mini-GBIC Module (J8706A)
- ProCurve Switch 5400zl 4-port 10-GbE X2 Module (J8707A)
- ProCurve Switch 5400zl 4-port 10-GbE CX4 Module (J8708A)

Each of the copper-based line interface modules provides integrated PoE capability.

24-port 10/100/1000

20-port 10/100/1000 + 4-port Mini-GBIC

24-port Mini-GBIC

4-port 10-GbE X2

4-port 10-GbE CX4

Figure 8. ProCurve Switch 5400zl Series Line Interface Modules

Power Supply Configurations

The 5406zl provides slots for two internal supplies and requires at least one internal power supply. An additional internal power supply may be added for 1+1 redundancy system power or to provide additional PoE power.

The 5412zl provides slots for four internal supplies and requires at least two internal power supplies. Two additional internal power supplies may be added for 1+1 system power redundancy and/or to provide additional PoE power.

1+1 power supply redundancy implies providing twice the minimum number of supplies required to power the system.

The table below shows the ProCurve Switch 5400zl series minimum and maximum internal power supply configurations, as well as internal combined power with external power shelf configurations.

Switch Model	Minimum Internal	Maximum Internal	Internal + External
5406zl	1	2	4
5412zl	2	4	6

Table 1. Minimum and Maximum Power Supply Support

When deciding on which and how many power supplies to configure for a 5400zl, the following criteria can be used to guide the decision:

- How much power will be required for each PoE port, full (15.4 watts) or phone (8 watts) power?
- How many network devices will require power (how many PoE ports are needed)?
- Is redundant power required (for internal power and/or PoE power)?

Note that the power supplies provide a "pool" of power for all line interface modules to draw from, that is, PoE power is not limited on a per line interface module basis.

Power Supply Configuration	Total PoE Power (Watts)	Number of PoE ports at 15.4 W		Number of PoE ports at 8 W	
		No Redundancy	With 1+1 Redundancy	No Redundancy	With 1+1 Redundancy
1-875W (J8712A)	273	17	N/A	34	N/A
2-875W (J8712A)	546	35	17	68	34
1-1500W (J8713A)	900	58	N/A	112	N/A
2-1500W (J8713A)	1800	116	58	144 *	112
4-1500W ** (J8713A)	3600	144 *	114	144 *	144*

Table 2. Examples of Maximum Number of PoE Ports for 5406zl

* 5406zl fully loaded with 6 modules provides up to 144 ports

** Using the ProCurve Switch zl Power Supply Shelf

The table above shows examples of the maximum number of PoE ports that can be supported by the 5406zl switch. One column shows the maximum number of PoE ports at full power (Class 0 – 15.4 watts) and the next column shows the maximum number of PoE ports at phone power (8 watts) for each configuration. These categories are further subdivided based on whether a second internal power supply is used to provide power redundancy in the event one of the power supplies fails.

Note that two 875W internal power supplies cannot supply sufficient PoE power to a 5406zl switch fully populated with 24-port line interface modules (144 ports). Two 1500W power supplies can be used to supply full PoE power (15.4 watts) to 116 ports.

Power Supply Configuration	Total PoE Power (Watts)	Number of PoE ports at 15.4 W		Number of PoE ports at 8 W	
		No Redundancy	With 1+1 Redundancy	No Redundancy	With 1+1 Redundancy
2-875W (J8712A)	546	35	17	68	34
4-875W (J8712A)	1092	70	35	136	68
2-1500W (J8713A)	1800	116	58	225	112
4-1500W (J8713A)	3600	233	116	288 *	225
6-1500W ** (J8713A)	5400	288 *	165	288 *	288 *

Table 3. Examples of Maximum Number of PoE Ports for 5412zl

* 5412zl fully loaded with 12 modules provides up to 288 ports

** Using the ProCurve Switch zl Power Supply Shelf

The table above shows examples of the maximum number of PoE ports that can be supported by the 5412zl switch. One column shows the maximum number of PoE ports at full power (Class 0 – 15.4 watts) and the next column shows the maximum number of PoE ports at phone power (8 watts) for each configuration. These categories are further subdivided based on whether the internal power supplies (1 of 2 or 2 of 4) are to provide power redundancy in the event a power supply fails.

Note that four 875W internal power supplies cannot supply sufficient PoE power to a 5412zl switch fully populated with 24-port line interface modules (288 ports). Four 1500W power supplies can be used to supply full PoE power (15.4 watts) to 233 ports.

For additional information about power supply configurations, see the *ProCurve Switch 5400zl/3500yl Series Ordering Guide* that can be found in the ProCurve Networking Reference Library at http://www.hp.com/rnd/library.

Specifications

Management Module

The ProCurve Switch 5406zl and 5412zl use a common management module that provides overall chassis management. The figure below illustrates the major components of the management module.

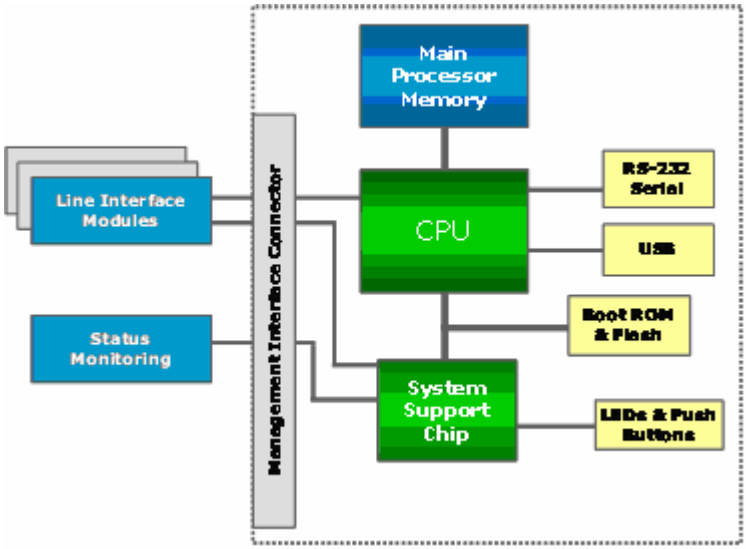

Figure 9. ProCurve Switch 5406zl Management Module Block Diagram

Processor

The CPU processor is a Freescale PowerPC 8540 operating at 667 MHz.

Memory

SDRAM

Synchronous Dynamic RAM is used for the storage of uncompressed executable code and data structures. The SDRAM consists of a 256 MB DDR-1 DIMM in the base module, expandable up to 1GB. The DDR-1 interface is 64-bits running at 166 MHz bus speed (333 MHz data rate).

FLASH

The flash consists of a 128 MB Compact Flash expandable up to 1GB, and a 4 MB mirror-bit flash. The mirror-bit flash is used for initial boot code. The Compact Flash is used for non-volatile configuration storage, and compresses code storage. The Compact Flash is socketed for future upgrade capability.

The Compact Flash may be programmed in a bulk fashion or one sector at a time. Since all application code is executed out of SDRAM, the Compact Flash may be programmed while the router is operational. The Compact Flash is sized such that a back-up copy of an older revision of application code may also be stored.

Console Port

The management module incorporates one RS-232 serial port for local management and configuration. This port uses a DB-9 male connector mounted on the front panel. To connect to the console, a standard null modem cable is used equivalent to that used for other ProCurve switches like the ProCurve 5300xl series.

Auxiliary Port

The management module includes a USB auxiliary port for use in future releases of the software.

Line Interface Modules

The following is a description of the line interface modules supported by the ProCurve Switch 5406zl and 5412zl.

ProCurve Switch 5400zl 24p 10/100/1000 PoE Module (J8702A)

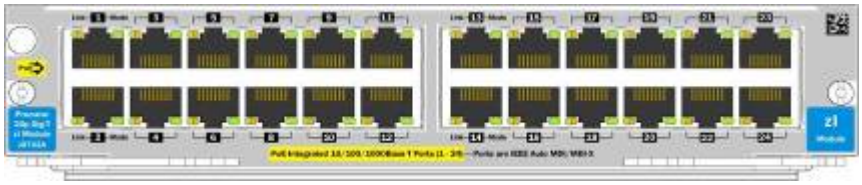

Description

This line interface module has 24 10/100/1000Base-T ports that provide gigabit-over-copper connectivity for wiring closets, enabling high-density gigabit connectivity to the desktop over Category 5 copper cabling. Each port is capable of providing IEEE 802.3af compliant Power over Ethernet (PoE) to power IP Phones, Wireless Access Points and other devices. Pre-standard powered devices can also be supported.

Ports

IEEE Auto-MDI: yes

Duplex: half or full

Connectors: RJ-45

ProCurve Switch 5400zl 20-port 10/100/1000 + 4-port Mini-GBIC Module (J8705A)

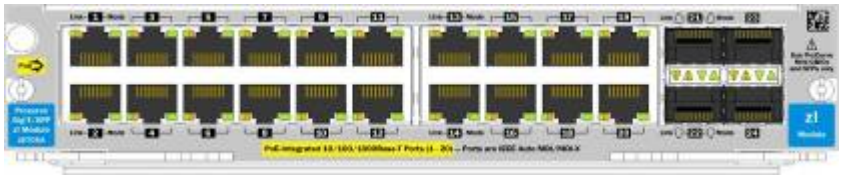

Description

This line interface module has 20 10/100/1000Base-T ports that provide gigabit-over-copper for wiring closets, enabling high-density gigabit connectivity to the desktop over Category 5 copper cabling. Each port is capable of providing IEEE 802.3af compliant Power over Ethernet (PoE) to power IP Phones, Wireless Access Points and other devices.

In addition, this module provides 4 mini-GBIC ports for uplinks and intra-building connections. They can be trunked to provide up to 4 Gigabits of connectivity. Two such modules can have their ports trunked across both modules to provide module redundancy as well.

Ports

4 open mini-GBIC slots

IEEE Auto-MDI: yes

Duplex: half or full

Connectors: RJ-45

Mini-GBICs supported (ordered separately)

- J4858B ProCurve Gigabit-SX-LC Mini-GBIC
- J4859B ProCurve Gigabit-LX-LC Mini-GBIC
- J4860B ProCurve Gigabit-LH-LC Mini-GBIC
- J8177B ProCurve Gigabit 1000Base-T Mini-GBIC

ProCurve Switch 5400zl 24-port Mini-GBIC Module (J8706A)

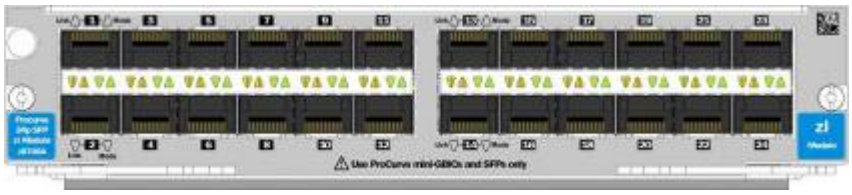

Description

This line interface module has 24 mini-GBIC ports and is appropriate for use as an aggregator in a distribution environment. This module supports the same mini-GBICs as the ProCurve Switch 5400zl 20-port 10/100/1000 + 4-port Mini-GBIC module previously described. The mini-GBICs are ordered separately.

Ports

24 open mini-GBIC slots

ProCurve Switch 5400zl 4-port 10-GbE X2 Module (J8707A)

Description

This line interface module has 4 10-GbE X2 Transceiver ports that support any combination of SR, LR, ER, or CX4 transceiver type. This provides maximum flexibility for connecting 10-GbE high speed downlinks to any other switch supporting that connection type. The wide variety of distances supported makes this module an ideal choice for intra-building connections. Ports can be trunked to provide higher throughput. Two such modules can have their ports trunked across both modules to provide module redundancy as well.

Ports

4 open X2 transceiver slots

Transceivers supported (ordered separately)

- J8436A ProCurve 10-GbE X2-SC SR Optic
- J8437A ProCurve 10-GbE X2-SC LR Optic
- J8438A ProCurve 10-GbE X2-SC ER Optic
- J8440A ProCurve 10-GbE X2-CX4 Transceiver

ProCurve Switch 5400zl 4-port 10-GbE CX4 Module (J8708A)

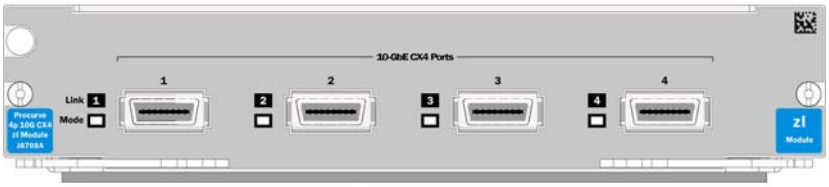

Description

This line interface module has 4 10-GbE CX4 ports.

Ports

4 10-Gigabit Ethernet ports (IEEE 802.3ak Type 10GBase-CX4)

Connectors: CX4

Transceivers supported (ordered separately)

- J8439A ProCurve 10-GbE CX4 Media Converter

Maximum Distance

15 m using CX4 cable

300 m using optical media converters and multimode fiber cable

Notes

Use CX4 10-GbE cable (0.5 m–15 m) or ProCurve 10-GbE CX4 Media Converter (J8439A)

ProCurve Switch 3500yl Series

The ProCurve Switch 3500yl series is another of most advanced intelligent edge switches in the ProCurve Networking product line. The 3500yl series includes 24-port and 48-port stackables. The foundation of these switches is a purpose-built, programmable ProVision ASIC that allows the most demanding networking features, such as QoS and security, to be implemented in a scalable yet granular fashion.

The ProCurve Switch 3500yl-24G-PWR supports 24 Gigabit interfaces and the ProCurve Switch 3500yl-48G-PWR supports 48 Gigabit interfaces. Four of the ports are dual-personality where each port can be used as either an RJ-45 10/100/1000 port (IEEE 802.3 Type 10Base-T; 802.3u Type 100Base-TX; 802.3ab 1000Base-T Gigabit Ethernet) or an open mini-GBIC slot (for use with mini-GBIC transceivers).

Both models provide integrated PoE on all 10/100/1000Base-T ports. Both models also have an expansion slot for an optional 4-port 10-GbE module. This line interface module is installed on the back panel of the switch.

The ProCurve Switch 3500yl series offers excellent investment protection, flexibility, and scalability, as well as ease of deployment, operation, and maintenance.

ProCurve Switch 3500yl-24G-PWR Intelligent Edge (J8692A)

ProCurve Switch 3500yl-48G-PWR Intelligent Edge (J8693A)

Figure 10. ProCurve Switch 3500yl Series

The table below shows examples of the maximum number of PoE ports that can be supported by the two 3500yl switch models. One column shows the maximum number of PoE ports at full power (Class 0 – 15.4 watts) and the other column shows the maximum number of PoE ports at typical phone power (8 watts).

For environments needing more PoE power, the ProCurve 620 Redundant/External Power Supply can be used. It doubles the available PoE power from 398W to 796W for up to two 3500 switches. The ProCurve 620 also provides redundant system power for up to two 3500 switches.

Switch Model	Total PoE Power (Watts)	Number of PoE ports at 15.4 W	Number of PoE ports at 8 W
3500yl-24G-PWR *	398	24	24
3500yl-48G-PWR *	398	24	46
3500yl-48G-PWR + ProCurve 620	796	48	48

Table 4. Examples of Maximum Number of PoE Ports for 3500yl Series

* Also number of ports that are 1+1 with the use of the ProCurve 620 Redundant/External Power Supply

Specifications

Processor

The CPU processor is a Freescale PowerPC 8540 operating at 667 MHz.

Memory

SDRAM

Synchronous Dynamic RAM is used for the storage of uncompressed executable code and data structures. The SDRAM consists of a 256 MB DDR-1 DIMM in the base module, expandable up to 1GB. The DDR-1 interface is 64-bits running at 166 MHz bus speed (333 MHz data rate).

FLASH

The flash consists of a 128 MB Compact Flash expandable up to 1GB, and a 4 MB mirror-bit flash. The mirror-bit flash is used for initial boot code. The Compact Flash is used for non-volatile configuration storage, and compresses code storage. The Compact Flash is socketed for future upgrade capability.

The Compact Flash may be programmed in a bulk fashion or one sector at a time. Since all application code is executed out of SDRAM, the Compact Flash may be programmed while the router is operational. The Compact Flash is sized such that a back-up copy of an older revision of application code may also be stored.

Console Port

An RS-232 serial port is supported for local management and configuration. The DB-9 serial port is located on the front panel of the ProCurve Switch 3500yl-24G-PWR and on the back panel of the ProCurve Switch 3500yl-48G-PWR. To connect to the console, a standard null modem cable is used that is equivalent to the cable used for other ProCurve switches like the ProCurve 5300xl series.

LED Status Indicators

Refer to the appendix for information about the LED status indicators of the ProCurve Switch 3500yl series.

Additional Line Interface Module

The ProCurve Switch 3500yl series supports one additional line interface module, the ProCurve Switch yl 10-GbE 2p CX4 + 2p X2 Module that is installed on the back panel of the switch.

ProCurve Switch yl 10-GbE 2p CX4 + 2p X2 Module (J8694A)

Description

This 10-GbE line interface module has two fixed CX4 ports and two X2 slots and is supported in both models of the ProCurve Switch 3500yl series and also the ProCurve Switch 6200yl-24G-mGBIC.

Ports

2 open X2 transceiver slots

2 10-Gigabit Ethernet ports (IEEE 802.3ak Type 10Gbase-CX4)

Duplex: full

Transceivers supported (ordered separately)

- J8437A ProCurve 10-GbE X2-SC LR Optic
- J8436A ProCurve 10-GbE X2-SC SR Optic
- J8438A ProCurve 10-GbE X2-SC ER Optic
- J8439A ProCurve 10-GbE CX4 Media Converter
- J8440A ProCurve 10-GbE X2-CX4 Transceiver

Notes

Only the two fixed CX4 ports on this module support the ProCurve 10-GbE CX4 Media Converter (J8439A).

ProCurve Switch 6200yl

The ProCurve Switch 6200yl-24G-mGBIC is an advanced Layer 3 stackable in 1U height. It has 24 mini-GBIC slots and an expansion slot for an optional 4-port 10-GbE module. Designed to be deployed as an aggregator of traffic from the edge to the core of the network, this switch supports a variety of Gigabit mini-GBICs, such as SX, LX, LH, and 1000Base-T.

The 6200yl switch has the same features as the 5400zl and 3500yl switches, but comes standard with Premium License features. Thus the routing protocols are already available as part of the aggregator switch use model.

Like the ProCurve Switch 5400zl and 3500yl series, the foundation for this switch is a purpose-built ProVision ASIC that allows the most demanding networking features, such as QoS and security, to be implemented in a scalable yet granular fashion.

With its high performance architecture, 10-GbE capability, and programmable ASIC, the ProCurve Switch 6200yl-24G-mGBIC offers excellent investment protection, flexibility, and scalability.

The ProCurve 620 Redundant/External Power Supply can be used to supply RPS power to the Switch 6200yl for high availability environments.

ProCurve Switch 6200yl-24G-mGBIC (J8992A)

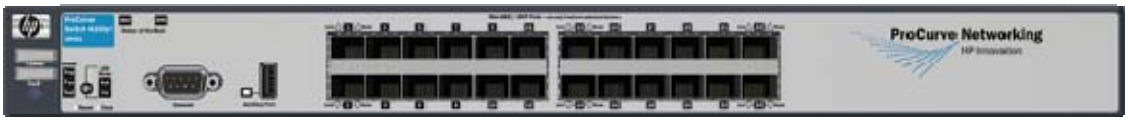

Figure 11. ProCurve Switch 6200yl

Processor

The CPU processor is a Freescale PowerPC 8540 operating at 667 MHz.

Memory

SDRAM

Synchronous Dynamic RAM is used for the storage of uncompressed executable code and data structures. The SDRAM consists of a 256 MB DDR-1 DIMM in the base module, expandable up to 1GB. The DDR-1 interface is 64-bits running at 166 MHz bus speed (333 MHz data rate).

Flash

The flash consists of a 128 MB Compact Flash expandable up to 1GB, and a 4 MB mirror-bit flash. The mirror-bit flash is used for initial boot code. The Compact Flash is used for non-volatile configuration storage, and compresses code storage. The Compact Flash is socketed for future upgrade capability.

The Compact Flash may be programmed in a bulk fashion or one sector at a time. Since all application code is executed out of SDRAM, the Compact Flash may be programmed while the router is operational. The Compact Flash is sized such that a back-up copy of an older revision of application code may also be stored.

Console Port

An RS-232 serial port is supported for local management and configuration. The DB-9 serial port is located on the front panel of the ProCurve Switch 6200yl-24G-mGBIC. To connect to the console, a standard null modem cable is used that is equivalent to the cable used for other ProCurve switches like the ProCurve 5300xl series.

LED Status Indicators

The front panel of the ProCurve Switch 6200yl-24G-mGBIC has the same LED status indicators as the ProCurve Switch 3500yl series. Refer to the appendix for information about the LED status indicators.

Additional Line Interface Module

The ProCurve Switch 6200yl-24G-mGBIC supports one additional line interface module, the ProCurve Switch yl 10-GbE 2p CX4 + 2p X2 Module. This module is installed on the back panel of the switch. This is the same module that is supported by the ProCurve Switch 3500yl series. Refer to the section covering the ProCurve Switch 3500yl series for details about this line interface module.

Overview of Features and Benefits

The ProCurve Switch 5400zl, 3500yl, and 6200yl series use the same software image base. For the ProCurve Switch 6200yl, the Premium License feature group is standard. For the ProCurve Switch 5400zl and 3500yl series, you have the choice of using the Intelligent Edge feature group or the Premium License feature group for an additional fee. The Premium License feature group supports additional Layer 3 services; PIM-SM, PIM-DM, OSPF-ECMP, and VRRP. The primary differences among these switch families are hardware related and include such aspects as port density and the number of power supplies and fans.

The following summary of features and benefits applies to the ProCurve Switch 5400zl, 3500yl, and 6200yl series. Any differences exist among the switches are noted.

Feature	Benefit
Performance	
ProVision ASIC technology	Powered by the ProVision ASICs, the switch families offer state of the art high capacity switch fabric performance—692 Gbps for the 5412zl, 346 Gbps for the 5406zl, 173 Gbps for the 3500yl-48G-PWR, and 115 Gbps for the 3500yl-24G-PWR and 6200yl-24G-mGBIC.
Selectable queue configurations	Increase performance by selecting the number of queues and associated memory buffer that best meet the requirements of network applications
Security Enabled:	
Virus throttle	Connection Rate Filtering thwarts viruses from spreading by blocking routing from certain hosts exhibiting abnormal traffic behavior.
ICMP rate-limiting	Throttles denial-of-service (DoS) attacks or other malicious behaviors that uses high volume ICMP traffic.
Switch CPU protection	Provides automatic protection against malicious network traffic trying to shut down the switch.
Detection of malicious attacks	Monitors 10 types of network traffic and send warning when anomoly that can potentially caused by malicious attacks is detected.
DHCP protection	Blocks DHCP packets from unauthorized DHCP servers, preventing denial-of-service attack
BPDU port protection	Blocks Bridge Protocol Data Unit (BPDU) on ports that do not require BPDU, preventing forged BPDU attack
Dynamic ARP protection	Blocks ARP broadcast from unauthorized hosts, preventing eavesdropping or data thief of network data
Dynamic IP lockdown (target availability 2Q 2007)	works with DHCP protection to block traffic from unauthorized host, preventing IP source address spoofing
Filtering capabilities	Include fast, flexible Access Control Lists (ACLs), up to 3,000 per module, source port, multicast MAC address, and other protocol-based filtering capabilities.
Identity Driven Manager	ProCurve Identity Driven Manager (IDM) to dynamically apply security, access and performance settings to infrastructure devices based on approved user, location and time.
Port security	Port security, MAC lockdown and MAC lockout protection for restricting access to the network through a switch port.
User Authentication	Client-based access control using 802.1X, Web-based, MAC-based authentication, RADIUS, and TACACS+. At initial release, the combinations of authentication methods allowed simultaneously on a port are 802.1X/Web and 802.1X/MAC.

Feature	Benefit
Management security	SSH v1 and v2, SSLv3 and TLSv1, TACACS+, and Secure FTP for management access. Management VLAN–Limit CLI/GUI/telnet/SNMP access to the switch to a particular VLAN. Manager authorized list limits CLI/Web/Telnet/SNMP management access to the switch to particular authorized hosts.
Redundancy and High Availability	
Redundancy Protocols	Protocols providing high availability include 802.1Q Multiple Spanning Tree Protocol, Switch Meshing, and Virtual Router Redundancy Protocol (Premium License).
Power Supplies	Optional redundant power supplies available for 5400zl series.
Management Module	Removable/replaceable management module for 5400zl series.
Operational Intelligence	
Policy Enforcement Engine	Policy Enforcement Engine is user configured to select packets that are then forwarded or dropped (based on ACLs, QoS, and Rate Limiting). The engine is fast, and can look for multiple variables, such as an IP address and port number, in a single pass through a packet. Provides a common user experience regardless of which switch the user is connected.
Operational Flexibility	
Versatile Intelligent Port	Versatile Intelligent Port supports flexible client port connections as every copper port is: 10/100/1000 capable, PoE enabled, all L3 services applied, and 8 hardware queues. All copper ports are created equal so any 10/100/1000 copper port can be connected to any 10/100/1000 copper client. The 5400zl and 3500yl switches support copper ports, whereas the 6200yl switch supports fiber only.
PoE Power	Power supply choice for optimum PoE power on 5400zl.
Port Density	Choice of 24 to 288 ports in single consistent product family.
L3 Services at L2	ACLs, QoS, and other features can use Layer 3 information such as IP addresses without requiring the switch to have routing feature set.
Intelligent Edge/ Premium License	Flexible approach to licensing. The 5400zl and 3500yl switches can run the Intelligent Edge feature group initially and then be upgraded later to run the Premium License feature group, if its advanced Layer 3 features are needed. The Layer 3 features include VRRP, PIM-SM, PIM-DM, and OSPF-ECMP. A Premium License can be transferred to another switch, as long as the license stays in the same hardware family. Premium License is standard on the 6200yl.
Convergence Ready	
Bandwidth Shaping / Control	Guaranteed minimums can be applied to traffic that must always get through the switch. Enforced maximums can be used to limit problem clients' bandwidth to no more than a set amount.
QoS	QoS based on TCP/UDP ports and other variables allows precise control of packet priority. Eight priority levels mapped to eight hardware queues can be used to set outgoing IP priority.
L2/L3 Jumbo Frames	Layer 2 / Layer 3 jumbo frames provide scalability in throughput.
Supports Standard Protocols	
IP Routing	Supports RIPv1, RIPv2, OSPF, OSPF-ECMP and static routes.
VLANs	Supports 802.1Q, 2048 concurrent VLANs (4094 in later release), Group VLAN Registration Protocol (GVRP).
IP Multicast	Supports PIM-SM, PIM-DM, and IGMP data-driven snooping.

Feature	Benefit
Low Cost of Ownership	
Future-proofing	Versatile Intelligent Port implies all Gigabit copper ports on the 5400zl and 3500yl series are PoE enabled. The 5400zl series chassis management module is upgradeable in order to meet future needs for processing power. Programmable ASICs allow some future requirements to be implemented without replacing the switch. Intelligent Edge features can be upgraded to Premium License features on the 5400zl and 3500yl series, allowing an edge switch to be redeployed as a distribution switch. The base power supply for the 5400zl series provides sufficient power for PoE on some ports. Power supplies can be added in order to provide PoE on all copper ports.
Lifetime warranty	Next-business-day advance replacement for as long as you own the product (available in most countries).
Lifetime SW Updates	No additional charge for software feature enhancements to the Intelligent Edge and Premium licenses.

Feature Set Summary

Additional details about the ProCurve Switch 5400zl, 3500yl, and 6200yl series features include the following:

Security Features

- Virus Throttle: Connection Rate Filtering thwarts virus spreading by blocking routing from certain hosts exhibiting abnormal traffic behavior.
- ICMP throttling: Defeats ICMP denial-of-service attacks by enabling any switch port to automatically throttle ICMP traffic.
- Filtering Capabilities: Includes fast, flexible Access Control Lists (ACLs), up to 3,000 per module (in later release, more precise detailed control via the fast Policy Enforcement Engine), source port, multicast MAC address, and other protocol-based filtering capabilities.
- Detection of malicious attacks: Monitors ten types of network traffic and sends a warning if an anomaly occurs signaling the detection of a potential malicious attacks..
- DHCP protection: Blocks DHCP packets from unauthorized DHCP servers, preventing denial-of-service attack.
- BPDU port protection: Blocks Bridge Protocol Data Unit (BPDU) on ports that do not require BPDU, preventing forged BPDU attack.
- Dynamic ARP protection: Blocks ARP broadcast from unauthorized hosts, preventing eavesdropping or data thief of network data.
- Dynamic IP lockdown: Works with DHCP protection to block traffic from unauthorized host, preventing IP source address spoofing.
- Multiple user authentication methods:
    - IEEE 802.1X: Industry-standard way of user authentication using an 802.1X supplicant on the client in conjunction with a RADIUS server
    - Web-based authentication: Similar to 802.1X, provides a browser-based environment to authenticate clients that do not support the 802.1X supplicant
    - MAC-based authentication: Client is authenticated with the RADIUS server based on the MAC address of the client; useful for clients that have minimal or no user interface
- Authentication flexibility:
    - Multiple 802.1X users per port: Provides authentication of multiple 802.1X users per port; prevents user “piggybacking” on another user's 802.1X authentication

    - Concurrent 802.1X and Web or MAC authentication schemes per port: At initial release, the combinations of authentication methods allowed simultaneously on a port are 802.1X/Web and 802.1X/MAC.
- Access control lists (ACLs): Provide filtering based on the IP field, source/destination IP address/subnet, and source/destination TCP/UDP port number on a per-VLAN or per-port basis.
- Identity-driven ACL: Enables implementation of a highly granular and flexible access security policy specific to each authenticated network user
- Port security: Prevents unauthorized access using MAC address lockdown
- MAC address lockout: Prevents configured particular MAC addresses from connecting to the network
- Source-port filtering: Allows only specified ports to communicate with each other
- Security banner: Displays customized security policy when users log in to the switch
- Management access:
    - All access methods—CLI, GUI, or MIB—are securely encrypted through SSHv2, SSL, and/or SNMPv3
    - RADIUS and TACACS+: Can require either RADIUS or TACACS+ authentication for secure switch CLI logon
    - Secure FTP: Allows secure file transfer to/from the switch and protects against unwanted file downloads or unauthorized copying of switch configuration file

QoS Functions

Layer 4 prioritization: Enables prioritization based on TCP/UDP ports.

Traffic prioritization: Allows real-time traffic classification into 8 priority levels mapped to 8 queues

Bandwidth shaping using:

- Rate limiting: Per-port ingress-based enforced bandwidth maximums.
- Guaranteed minimums: Per-port, per-queue egress-based guaranteed bandwidth minimums.

Class of Service (CoS): Sets 802.1p priority tag based on IP address, IP Type of Service (ToS), L3 protocol, TCP/UDP port number, source port, and DiffServ.

Convergence

- IP multicast routing: Includes PIM Sparse and Dense modes to route IP multicast traffic.
- IP multicast data-driven IGMP: Automatically prevents flooding of IP multicast traffic.
- LLDP-MED (Media Endpoint Discovery): A standard extension of LLDP that stores values for parameters such as QoS and VLAN to automatically configure network devices such as IP phones.
- iSCSI support: Enables the deployment of Ethernet storage area network solutions using the iSCSI standard.

Layer 2 switching

- ProCurve switch meshing: Dynamically load balances across multiple active redundant links to increase available aggregate bandwidth.
- VLAN support and tagging: Support for complete 802.1Q standard and 2,048 VLANs simultaneously.
- 802.1v protocol VLANs: Isolate select non-IPv4 protocols automatically into their own VLANs.
- GVRP: Group VLAN Registration Protocol allows automatic learning and dynamic assignment of VLANs.

Bridging Protocols

- MSTP: Provides high link availability in multiple VLAN environments by allowing multiple spanning trees; encompasses 802.1D STP and 802.1w RSTP.

Routing Protocols

- RIP (v1, v2, and v1 compatible v2)
- OSPF with ECMP (Equal Cost Multi-Path) and NSSA
- Static IP routes: To manually add routes directly to the routing table
- 10,000 network address routes, 65,536 L3 host address routes
- UDP helper function: UDP broadcasts can be directed across router interfaces to specific IP unicast or subnet broadcast addresses and prevent server spoofing for UDP services such as DHCP
- Loopback interface address: Define an interface in RIP and OSPF that can always be reachable, improving diagnostic capability.
- IPv4 routing
- IPv6-ready in hardware

Multicast Protocols

- IGMP data-driven
- PIM-SM, PIM-DM

High Availability and Redundancy

- VRRP: Virtual Router Redundancy Protocol (requires Premium License) allows groups of two routers to dynamically back each other up to create highly available routed environments
- 802.3ad LACP: Link Aggregation Control Protocol and ProCurve trunking support up to 36 trunks, each with up to 8 links (ports) per trunk.
- Port trunks, or link aggregation groups, can operate across multiple modules to increase redundancy.
- Supports various redundant power supply configurations:
    - The 5406zl provides slots for two internal supplies. This can provide 1+1 redundancy for both PoE and system power needs.
    - The 5412zl provides slots for four internal supplies. This can provide 1+1 redundancy similar to the 5406zl. You can use the internal supplies for 1+1 backup; two backing up two for both PoE and system power needs.
- Hot-swappable/hot-insertable line interface modules, so as not to interrupt network operation.
- Management module is removable/upgradeable: The switch does not have to be powered off to remove a management module from a 5400zl series. However, when the management module is removed all ports will lose communication and the system will be powered down.
- Other Layer 2 and Layer 3 redundant protocols include MSTP, ProCurve switch meshing, and OSPF-ECMP.

Management

- Remote intelligent mirroring: Mirrors ingress/egress ACL-selected traffic from a switch port or VLAN to a local or remote switch port anywhere on the network.
- RMON, XRMON, and sFlow: Provide advanced monitoring and reporting capabilities for statistics, history, alarms, and events.
- Command authorization: Leverages RADIUS to link a custom list of CLI commands to individual network administrator's login; also provides audit trail
- LLDP: Link Layer Discovery Protocol allows automated device discovery for easy mapping by network management applications.
- Friendly port names: Allow assignment of descriptive names to ports.
- Dual flash images: Provide independent primary and secondary OS files for backup while upgrading.
- Multiple configuration files: Allows multiple configuration files to be stored to flash image.

- Uni-directional Link Detection (UDLD): Monitors a link between two switches and blocks the ports on both ends of the link if the link goes down at any point between the two devices
- USB support: Allows users to copy files from the switch to a USB device.
- Configuration, monitoring and management tasks can be performed using the CLI (through direct console or Telnet) connection, Web browser, Menu interface, and ProCurve Manager.

Standards and Protocols

RFC 768 UDP	RFC 783 TFTP
RFC 792 ICMP	RFC 793 TCP
RFC 826 ARP	RFC 854 Telnet
RFC 951 BootP	RFC 1058 RIP
RFC 1213 MIB II	RFC 1492 TACACS+
RFC 1493 Bridge MIB	RFC 1724 RIPv2 MIB
RFC 1542 BootP	RFC 1850 OSPF MIB
RFC 2021 RMON v2 probe configuration	RFC 2030 Simple Network Time Protocol
RFC 2096 IP Forwarding Table MIB	RFC 2138 RADIUS
RFC 2328 OSPFv2 [1]	RFC 2453 RIPv2
RFC 2474 DiffServ Precedence	RFC 2597 DiffServ Expedited Forwarding
RFC 2598 DiffServ Assured Forwarding	RFC 2613 SMON
RFC 2618 RADIUS Client MIB	RFC 2620 RADIUS Accounting MIB
RFC 2665 Ethernet MIB	RFC 2668 802.3 MAU MIB
RFC 2674 802.1p and IEEE 802.1Q Bridge MIB	RFC 2737 Entity MIB
RFC 2787 VRRP MIB	RFC 2819 RMON MIB [2]
RFC 2863 Evolution of Interface	RFC 2866 RADIUS accounting
RFC 3046 DHCP Relay Agent Information Option	RFC 3376 IGMPv1/v2/v3
RFC 3768 VRRP [1]	IEEE 802.1AB Link Layer Discovery Protocol
IEEE 802.1D Spanning Tree	IEEE 802.1p Priority
IEEE 802.1Q GVRP	IEEE 802.1Q VLAN tagging
IEEE 802.1s Multiple Spanning Tree	IEEE 802.1w Rapid Convergence Spanning Tree
IEEE 802.1X Network Login	IEEE 802.3ad Link Aggregation Control Protocol
IEEE 802.3af Power over Ethernet	IEEE 802.3x Flow Control
ANSI/TIA-1057 LLDP Media Endpoint Discovery	DHCP Relay
HTML and Telnet management	IGMPv3
PIM Dense Mode [1]	PIM Sparse Mode [1]
Secure Sockets Layer	sFlow
SNMPv1/v2c/v3	SSHv1/SSHv2 Secure Shell
XRMON	RFC 2925 Ping MIB
RFC 3101 OSPF NSSA [1]	

Table 5. Supported Standards and Protocols

Notes:

1 Requires Premium License

2 RFC 2819 RMON MIB four groups: 1 (statistics), 2 (history), 3 (alarm) and 9 (events)

Performance and Capacity

Capacity and Performance Features Comparison

The following table provides a comparison of several categories of capacity and performance levels for the ProCurve Switch 5400zl, 3500yl and 6200yl series. These include the routing and switching capacity (Gbps), switch fabric speed (Gbps), maximum number of 1 Gbps that can operate concurrently at wire-speed, maximum number of 10 Gbps ports that can operate concurrently at wire-speed, and the size of the routing table (entries).

Capacity / Performance Feature	ProCurve Switch 5406zl	ProCurve Switch 5412zl	ProCurve Switch 3500yl-24G-PWR	ProCurve Switch 3500yl-48G-PWR	ProCurve Switch 6200yl-24G-mGBIC
Routing / Switching capacity (Gbps)	288	576	101	148	101
Switch Fabric Speed (Gbps)	346	692	115	173	115
Max 1 Gigabit ports at wire-speed	144	288	24	48	24
Max 10 Gigabit ports at wire-speed	12	24	2	2	2
Routing table size (entries)	10,000	10,000	10,000	10,000	10,000

Table 6. Capacity and Performance Features Comparison of 5400zl, 3500yl and 6200yl Series

Per-port Buffer Sizes

The ProCurve Switch 5400zl, 3500yl and 6200yl series have eight hardware queues to which prioritized traffic can be assigned. Software can designate certain amount of output memory buffer for each queue to minimize impact from events such as bursty network traffic that exceed line rate. Users can choose to configure the number of queues --- 2, 4, or 8 --- so that they can optimize the amount of memory buffer per queue for their applications. For example, if a user has only 2 priorities on the network, that user should configure for 2 queues so that all the buffer memory is assigned to only 2 queues. The table below identifies the buffer sizes allocated on a per-port basis to the normal and other priority queues in each configuration.

Port Speed	Size of Normal Priority Queue (KB)	Size of Other Priority Queues (KB)
1 Gbps	194 KB (30%)	70 KB (10%)
10 Gbps	225 KB (30%)	77 KB (10%)

Table 7a. Per-port Priority Queue Buffer Sizes for 8 queues

Port Speed	Size of Normal Priority Queue (KB)	Size of Other Priority Queues (KB)
1 Gbps	480 KB (70%)	68 KB (10%)
10 Gbps	538 KB (70%)	76 KB (10%)

Table 7b. Per-port Priority Queue Buffer Sizes for 4 queues

Port Speed	Size of Normal Priority Queue (KB)	Size of Other Priority Queues (KB)
1 Gbps	618 KB (90%)	68 KB (10%)
10 Gbps	691 KB (90%)	76 KB (10%)

Table 7c. Per-port Priority Queue Buffer Sizes for 2 queues

Routing and Forwarding Tables

The ProCurve Switch 5400zl, 3500yl, and 6200yl series have several routing and forwarding table features that enhance routing and switching performance.

Every line interface module has its own “best match prefix” routing table that contains IP routes and is used for determining how to route the vast majority of incoming packets. Using the best match prefix routing table is extremely fast and enables wire-speed routing to be achieved.

If the best match prefix routing table does not contain an entry that can be used to determine the route of a received packet, then the main routing table is used. The main routing table can contain up to 10,000 routing table entries.

For Layer 2 forwarding, each switch has a forwarding or MAC table that can contain up to 65,536 entries.

Optimizing the 10-GbE Port Configuration

The 10-GbE modules (J8707A, J8708A, and J8694A) used in the ProCurve Switch 3500yl, 5400zl, and 6200yl series are designed to deliver full 10 Gbps wire-speed to each port, where either one or two ports are in a linked state with another device. When three or four 10-GbE ports are in a linked state, the module supports an aggregate bandwidth of 28.8 Gbps across the linked ports.

As illustrated in the figure below, internally, there are two 14.4-Gbps channels between each 4-port 10-GbE module and the switch fabric.

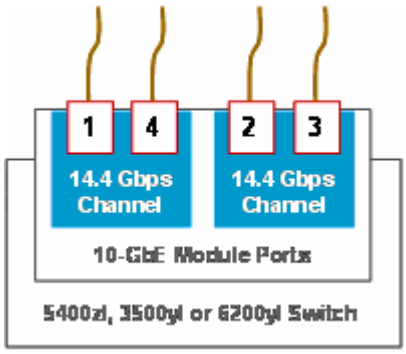

Figure 12. 10-GbE Module Architecture Showing Ports Grouped by 14.4 Gbps Channel

- When any two 10-GbE ports on a given module are in a linked state, each port automatically operates on its own channel, which guarantees 10 Gbps of bandwidth for each port.
- When more than two ports of a 10-GbE module are in a linked state, ports 1 and 4 are statically mapped to share one 14.4 Gbps channel, while ports 2 and 3 are statically mapped to share the other 14.4 Gbps channel.
    - If only one port in a given channel is in a linked state, then that port operates at wire speed and the other port uses no bandwidth.
    - If both ports in a given channel are in a linked state, then the 14.4 Gbps of bandwidth is balanced fairly between the two ports

For example, in an application where three ports on the module are needed and it is important to ensure that port 1 always has a full 10 Gbps available, the connection choices shown in the next figure would be used.

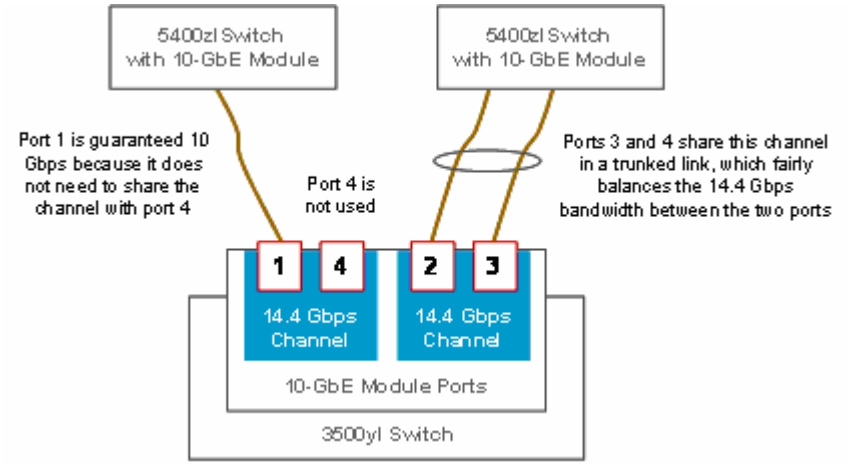
5400zl Switch with 10-GbE Module
5400zl Switch with 10-GbE Module
Port 1 is guaranteed 10 Gbps because it does not need to share the channel with port 4
Port 4 is not used
Ports 3 and 4 share this channel in a trunked link, which fairly balances the 14.4 Gbps bandwidth between the two ports
1
4
2
3
14.4 Gbps Channel
14.4 Gbps Channel
10-GbE Module Ports
3500yl Switch

Figure 13. Approach for Guaranteeing 10 Gbps on a Specific Port

Connection choices are also important where equally balanced bandwidth is needed, such as in a Virtual Router Redundancy Protocol (VRRP) application. This scenario is illustrated in the next figure.

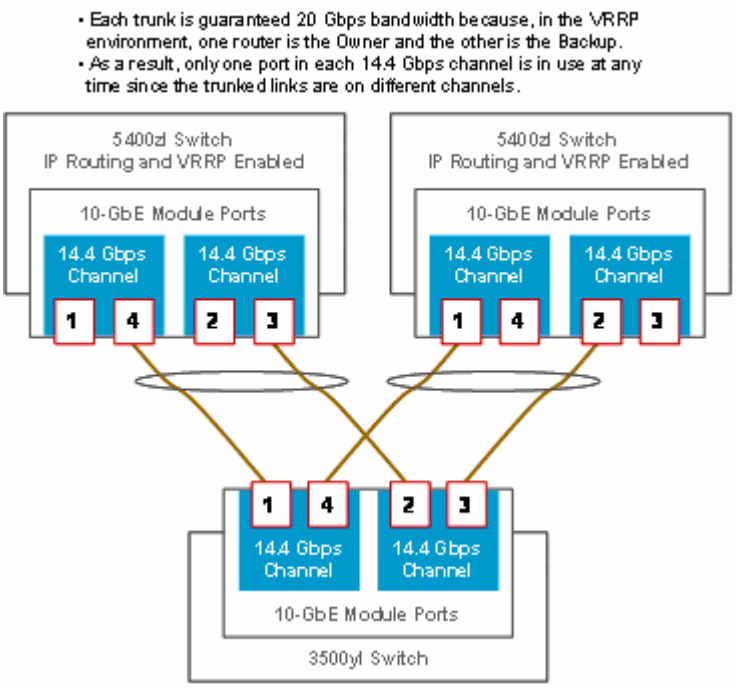
• Each trunk is guaranteed 20 Gbps bandwidth because, in the VRRP environment, one router is the Owner and the other is the Backup.
• As a result, only one port in each 14.4 Gbps channel is in use at any time since the trunked links are on different channels.
5400zl Switch IP Routing and VRRP Enabled
10-GbE Module Ports
14.4 Gbps Channel
14.4 Gbps Channel
1
4
2
3
5400zl Switch IP Routing and VRRP Enabled
10-GbE Module Ports
14.4 Gbps Channel
14.4 Gbps Channel
1
4
2
3
1
4
2
3
14.4 Gbps Channel
14.4 Gbps Channel
10-GbE Module Ports
3500yl Switch

Figure 14. Guaranteeing Equally Balanced Bandwidth in a VRRP Environment

Throughput and Latency Performance Data

The following three tables describe the performance capabilities of the 10 Gbps, 1 Gbps and 100 Mbps line interface module ports of the ProCurve Switch 5400zl, 3500yl and 6200yl series.

Over a range of packet sizes from 64 bytes to 1518 bytes, the performance levels achieved are described in terms of the following metrics:

- Throughput level percentage achieved when performing Layer 2 switching and Layer 3 routing; 100% in all cases.
- Throughput in the form of the number of packets per second when performing Layer 2 switching and Layer 3 routing.
- Average per packet latency in microseconds when performing Layer 2 switching and Layer 3 routing. These values are shown are listed for both First In First Out (FIFO) and Last In First Out (LIFO) processing. The LIFO values represent the packet transmit time, whereas the FIFO values include the switch decision time and the packet transmit time.
- Percentage of packets dropped in a full mesh configuration; 0% in all cases.

The following measurements were performed by ProCurve using test equipment manufactured by Ixia Communications (www.ixiacom.com). In these tests, the maximum number of supported ports on the ProCurve Switch 5412zl was used. Proportional performance results at 100% of throughput are achieved on the ProCurve Switch 5406zl, 3500yl-24G-PWR, 3500yl-48G-PWR, and 6200yl-24G-mGBIC for the maximum number of ports supported by each of those models. The 100 Mbps table is not applicable to the ProCurve Switch 6200yl-24G-mGBIC, since only 1 Gigabit and 10 Gigabit ports are supported.

10 Gbps	Throughput%	Packets per Second	Full Mesh % Drops	Latency (FIFO)		Latency (LIFO)	
Packet Size	L2/L3	L2/L3		L2 (µsecs)	L3 (µsecs)	L2 (µsecs)	L3 (µsecs)
64 [1]	100	14880952	0	1.95	2.05	1.9	2.0
128	100	8445946	0	2.01	2.11	1.9	2.0
256	100	4528986	0	2.31	2.31	2.1	2.1
512	100	2349624	0	2.62	2.72	2.2	2.3
1024	100	1197318	0	3.73	3.73	2.9	2.9
1280	100	961538	0	4.13	4.23	3.1	3.2
1518	100	811688	0	4.52	4.52	3.3	3.3

1 Gbps	Throughput%	Packets per Second	Full Mesh % Drops	Latency (FIFO)		Latency (LIFO)	
Packet Size	L2/L3	L2/L3		L2 (µsecs)	L3 (µsecs)	L2/L3	L2/L3
64	100	1488095	0	3.28	3.19	2.7	2.7
128	100	844595	0	3.79	3.79	2.7	2.7
256	100	452899	0	4.91	4.91	2.8	2.8
512	100	234962	0	7.16	7.16	3.0	3.0
1024	100	119732	0	11.76	11.76	3.5	3.5
1280	100	96154	0	14.00	14.2	3.7	3.9
1518	100	81274	0	16.20	16.20	4.0	4.0

100 Mbps	Throughput %	Packets per Second	Full Mesh % Drops	Latency (FIFO)		Latency (LIFO)	
Packet Size	L2/L3	L2/L3		L2 (µsecs)	L3 (µsecs)	L2/L3	L2/L3
64	100	148810	0	3.28	3.19	2.7	2.7
128	100	84459	0	3.79	3.79	2.7	2.7
256	100	45290	0	4.91	4.91	2.8	2.8
512	100	23496	0	7.16	7.16	3.0	3.0
1024	100	11973	0	11.76	11.76	3.5	3.5
1280	100	9615	0	14.00	14.2	3.7	3.9
1518	100	8127	0	16.20	16.20	4.0	4.0

Table 8. Throughput Performance

1 See the explanation about 10 Gigabit Performance Traffic Patterns in the next section.

10 Gigabit Performance Traffic Patterns

In the prior table, the performance levels for **10 Gigabit** ports assume the underlying traffic patterns reflect either one of the following minimum conditions to achieve wire-speed throughput.

- A single source traffic stream with an average packet size of 88 bytes or larger.
- Two or more source traffic streams of any packet size down to the minimum value of 64 bytes.

In the unlikely case where the average packet size is consistently smaller, the throughput will be less than wire-speed. For example, consider a worst case scenario where the average packet size is 64 bytes. This would result in a throughput of approximately 70% to 80% of the rated wire-speed capacity. ProCurve considers such minimum-sized, packet traffic scenarios being realized over an extended period of time to be extremely atypical and unlikely to be experienced by customers in the field.

Note that the limits described above do not apply to Gigabit ports.

Throughput Test

A fully meshed performance test sends packets from each port to every other port during the test. This type of test exercises both the modules and the backplane. These tests show the ProCurve Switch 5400zl, 3500yl and 6200yl series to be wire-speed on all ports simultaneously.

Latency Measurements

Latency is commonly measured as the amount of time it takes for a byte inside a packet to enter and then leave the switch. Latency statistics are typically documented as including both the processing time of the switch as it makes its forwarding decision and the time for the packet itself to enter and leave the switch. In the prior tables, this definition of latency corresponds to the FIFO latency statistics. The LIFO latency statistics that are also listed in the tables represent only the packet transmission time.

Almost all switches currently on the market are store and forward, so the entire packet is received into the switch before the switch begins to transmit the packet out the egress port. Including the packet receive time in the FIFO latency statistics is appropriate since this extra time is a contributing component of the overall transit time of the packet as it moves through the network.

The latency figures for the ProCurve Switch 5400zl series are consistently low. Latencies this low will not be a factor in general network operation, even with streaming video or VoIP applications. The LIFO latency values are fairly consistent across all packet sizes due to the fact that ingress and egress packet processors operate on the header of the frame (not the whole frame) while the full frame is buffered in and out of packet buffer memory. Memory transfers are scheduled to fit a full 1518-byte frame so frames are transferred in and out of memory in approximately the same amount of time regardless of packet size. While the frame headers are

being looked up and actions required for the frame on egress are being coordinated among interface modules, the frame is transferred through the switching fabric module.

Comparison of 5400zl with 5300xl and 8100fl

The table below provides a general comparison of the ProCurve Switch 5400zl and 3500yl series with the ProCurve Switch 5300xl and 8100fl series products. The features compared include the number of 1 Gigabit and 10 Gigabit ports supported, and routing table sizes.

The 8100fl series offers high performance, high availability, cost effective connectivity for intelligent edge devices while delivering a flexible, scalable, high port-density 1 Gigabit and 10 Gigabit Ethernet core networking solution. The 8100fl series offers two chassis configurations, the ProCurve Switch 8108fl and ProCurve Switch 8116fl. The ProCurve Switch 8108fl is an eight-slot chassis-based routing switch with support for up to eight 10GbE ports or 80 100/Gigabit Ethernet ports. The ProCurve Switch 8116fl is a 16-slot chassis-based routing switch with support for up to 16 10GbE or 160 100/Gigabit Ethernet ports.

The ProCurve Switch 5300xl series offers scalable layer 2, 3 and 4 switching in compact 4- and 8-slot modular form factors. The ProCurve Switch 5304xl is a four-slot chassis-based routing switch with support for up to 64 10/100/1000 Ethernet ports. The ProCurve Switch 5308xl is an eight-slot chassis-based routing switch with support for up to 128 10/100/1000 Ethernet ports.

Feature	ProCurve Switch 5400zl	ProCurve Switch 3500yl	ProCurve Switch 5300xl	ProCurve Switch 8100fl
Routing / Switching capacity (Gbps)	576	148	76	320
Chassis Height	7U (5406zl) 12U (5412zl)	1U	3U (5304xl) 5U (5308xl)	13U (8108fl) 20U (8116fl)
Maximum Number of 1 Gb Ports	288	48	128	160
Maximum Number of 10 Gb Ports	48	4	0	16
Power over Ethernet	Yes	Yes	Yes	No
Jumbo Frames	Yes	Yes	No	Yes
Number of Priority Queues	8	8	4	5
Number of ACLs (entries)	3,000	3,000	1000	1000
Number of VLANs	2048	2048	256	2048
Forwarding Table Size (entries)	64K	64K	16K	100,000
Routing Table Size (entries)	10,000	10,000	10,000	100,000
RIP v1/v2	Yes	Yes	Yes	Yes
OSPF / OSPF-ECMP	Yes	Yes	Yes	Yes
PIM-SM	Yes	Yes	No	No
PIM-DM	Yes	Yes	Yes	No
Spanning Tree	MSTP	MSTP	RSTP, MSTP	STP, RSTP, MSTP

Feature	ProCurve Switch 5400zl	ProCurve Switch 3500yl	ProCurve Switch 5300xl	ProCurve Switch 8100fl
LACP	Yes	Yes	Yes	Yes
Switch Redundancy	VRRP	VRRP	XRRP	VRRP
Redundant Power	Yes	Yes (external unit)	Yes	Yes
Redundant Switch Fabric	No	No	No	Yes
Redundant Management Module	No	No	No	Yes

Table 9. Comparison of 5400zl, 3500yl, 5300xl and 8100fl Series

The architecture used for ACL storage and processing is different among the ProCurve switches listed in the table above. The ProCurve Switch 5400zl and 3500yl series have the additional performance advantage gained through the use of the TCAM (see the appendix for more details). The ProCurve Switch 5400zl and 3500yl series also have an additional performance advantage in the processing of the routing table through the use of a “best match prefix” table (described previously in this section). The ProCurve Switch 5400zl and 3500yl series support the use of PIM-SM and PIM-DM, but not concurrently.

Services & Support

Lifetime Software Updates (Best in the Industry)

As with other ProCurve products, the ProCurve Switch 5400zl, 3500yl, and 6200yl series come with lifetime software updates. For as long as software updates are available for these switches, customers can download updates from the ProCurve web site free of charge.

Lifetime Warranty (Best in the Industry)

Warranties, and the ease of obtaining warranty service for the customer, is a product benefit that is easily overlooked in a technical evaluation, but ranks high as a concern of customers as they get ready to actually make a purchase decision. The ProCurve Switch 5400zl, 3500yl, and 6200yl series have a limited lifetime warranty, for as long as you own the product. If any part of the switch fails due to a defect in material or workmanship, including the power supply or fans, it will be replaced. In most parts of the world, the replacement unit is sent with next business day delivery in advance of the failing unit being returned to ProCurve.

Advance replacement gets the unit to the customers as fast as possible and minimizes down time by allowing the impaired unit to continue to be used if possible until the replacement unit arrives. This also allows for easy scheduling for when the actual unit swap occurs on the network. An optional upgrade to on-site replacement is also available in most parts of the world. Refer to the warranty statement that ships with the product for more details on warranty coverage.

The ProCurve Switch 5400zl, 3500yl, and 6200yl series warranties are industry-leading.

Free Telephone Support

ProCurve provides free pre-sales and post-sales telephone support during normal business hours to end users and ProCurve resellers through the ProCurve Customer Care Centers located world-wide.

Optional Support Services

In addition to free support services such as the warranty and telephone support, ProCurve offers an extensive range of fee-based support services to meet more specialized needs. The following optional services are available for the ProCurve Switch 5400zl, 3500yl, and 6200yl series:

- Onsite next business day
- Onsite in 4 hours same business day
- Onsite in 4 hours 24x7
- 6 hour call-to-repair
- 24x7 telephone support

Hewlett-Packard can also provide more broad-based services such as site surveys, installation services, and actual management of the network, depending on customer needs.

More information can be found at http://www.hp.com/rnd/services/index.htm or by contacting a local ProCurve sales office.

For more information on the ProCurve Networking Adaptive EDGE Architecture, see the ProCurve Networking website at http://www.hp.com/rnd/architecture/index.htm.

Appendix A: Premium License

This section of the reviewer's guide provides additional background information about the ProCurve Switch 5400zl, 3500yl, and 6200yl series products.

Intelligent Edge and Premium License

Introduced with the ProCurve Switch 5400zl and 3500yl series is the option of implementing *Intelligent Edge* or *Premium License*. For the ProCurve Switch 6200yl, *Premium License* is standard due to its recommended deployment at the distribution level as an aggregator of traffic from the edge to the core.

In the case of the 5400zl and 3500yl switches, Premium License implies the switch can run advanced Layer 3 features that include OSPF-ECMP, PIM-SM, PIM-DM, and VRRP. All the product features are included in the firmware image installed on the switch, but Premium License features require a license for those features to be enabled.

For ProCurve, moving to Premium License creates a distinction between the terms "Upgrade" and "Update". Upgrade means licensing Premium License features which add significant new functionality to the switch for an additional fee. Update means installing the latest version of firmware such as bug fixes or additional features for the switch that are available for free for as long as the customer owns the switch. A Premium License will enable Premium License updates.

ProCurve has taken a flexible approach with Premium License. For example, in the 5400zl and 3500yl series, all switches can run the Intelligent Edge feature group initially. A switch can later be upgraded to run the Premium License feature group, if its advanced features are needed. Also, if network configurations change, a Premium License can be transferred to another switch, as long as the license stays in the same hardware family.

There are two ways you can manage Premium Licenses:

- **Manually** — You use the My ProCurve portal and the switch CLI. This is a good method to use if you are managing a small number of licenses and do not already use the ProCurve Manager network management software.
- **PCM** — You use ProCurve Manager to perform the most common licensing functions. If you are already using PCM, this is more convenient than the manual method. There is a free copy of ProCurve Manager on the CD-ROM that accompanies your premium license.

For the manual method, you will require manger privilege level access on the switches to perform several CLI commands related to licensing. For the PCM method, PCM requires configured access to the devices.

Two methods for managing Premium Edge licensing on ProCurve Switch 5400zl and 3500yl series

- Manually using ProCurve web portal and CLI commands
  - Good for small number of licenses, requires manager level access
- ProCurve Manager (PCM)
  - Convenient for may licenses, handles most aspects
  - Requires PCM configured for access to each device and My ProCurve account information

Either method requires that you be registered on the *My ProCurve* web portal

Task	Manual Method	Using PCM
Install a license	Yes	Yes
Uninstall a license	Yes	Yes
View available registration IDs	Yes	No *
View license history	Yes	No *
Export license history	Yes	No *

* Can access via My ProCurve web portal

To use ProCurve Manager to install a Premium License, you first select the 5400zl or 3500yl switch in the PCM navigation tree and right-click it. Then, you select *Config Manager* > *License*

Software on the pop-up menus. The easy-to-use License Configuration Wizard window will then be displayed so that you can start the license installation process.

Note that before attempting to install a Premium License using ProCurve Manager, ProCurve Manager must be configured for device access to the 5400zl or 3500yl switch and your "My ProCurve" account (member ID and password) information.

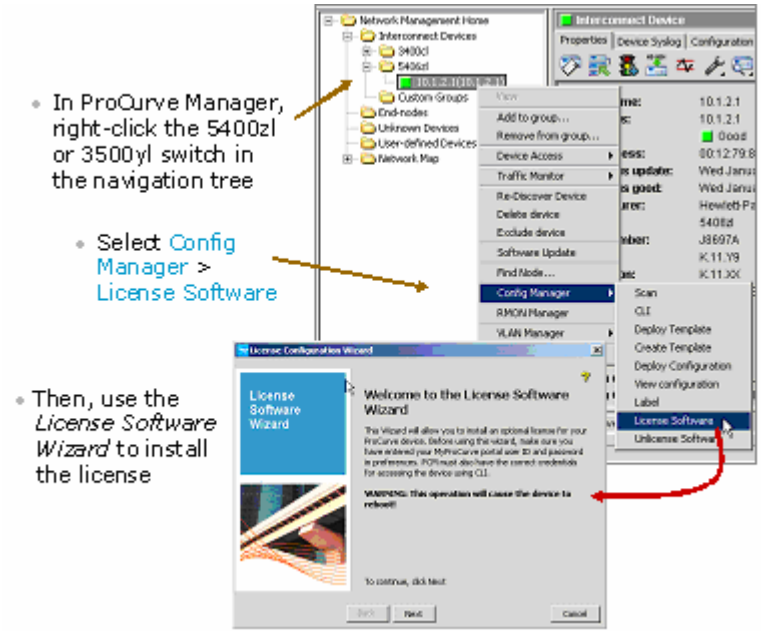
In ProCurve Manager, right-click the 5400zl or 3500yl switch in the navigation tree
Select Config Manager > License Software
Then, use the *License Software Wizard* to install the license
License Software Wizard
Welcome to the License Software Wizard

On a subsequent window, you select *premium-edge* for the Feature Group and type the registration ID you received with the purchase of the Premium License. The registration ID is found on a plastic laminated card that is shipped with the ProCurve switch.

After agreeing to the terms of the License Agreement, the wizard will display the status of the license installation process. As part of this process the following events occur:

- First, the Hardware ID is retrieved from the switch.
- Secondly, the license key is retrieved from the My ProCurve web portal.
- Lastly, the Premium License feature group is activated and the license is installed.

Appendix B: Policy Enforcement Engine

The ProVision ASIC architecture used in the ProCurve Switch 5400zl, 3500yl, 6200yl series brings a number of advanced capabilities to the network that offer a highly reliable, robust environment that leads to increased network uptime, keeping overall network costs down. One major feature is the ProVision *Policy Enforcement Engine* which is implemented in the ProVision ASIC of each line interface module.

Policy Enforcement Engine Benefits

The Policy Enforcement Engine has several benefits:

- **Granular policy enforcement**
  The initial software release on these products takes advantage of a subset of the full Policy Enforcement Engine capabilities, which will provide a common front end for the user interface to ACLs, QoS, Rate-Limiting, and Guaranteed Minimum Bandwidth controls. Fully implemented in later software releases, the Policy Enforcement Engine provides a powerful, flexible method for controlling the network environment. For example, traffic from a specific application (TCP/UDP port) can be raised in priority (QoS) for some users (IP address), blocked (ACL) for some other users, and limited in bandwidth (Rate-Limiting) for yet other users.

  The Policy Enforcement Engine provides fast packet classification to be applied to ACLs

and QoS rules, and Rate Limiting and Guaranteed Minimum Bandwidth counters. Parameters that can be used include source and destination IP addresses, which can follow specific users, and TCP/UDP port numbers and ranges, which are useful for applications that use fixed port numbers. Over 14 different variables can be used to specify the packets to which ACL, QoS, Rate Limiting, and Guaranteed Minimum Bandwidth controls are to be applied.

- **Hardware-based performance**
  As mentioned above, the Policy Enforcement Engine is a part of the ProVision ASIC. The packet selection is done by hardware at wire speed except in some very involved rules situations. Therefore, very sophisticated control can be implemented without adversely affecting performance of the network.

- **Works with Identity Driven Manager**
  Identity Driven Manager (IDM) provides the centralized command from the center to define the user policies to be used with each user. The IDM policy requests sent down to the switch are used to set up the user profile in the Policy Enforcement Engine so that the per-user ACL, QoS, and Rate-Limiting parameters can be used from the actual policy defined in IDM.

Wire-speed Performance for ACLs

At the heart of the Policy Enforcement Engine is a memory area called the Ternary Content Addressable Memory (TCAM) that is contained within the ProVision ASIC along with the surrounding code for the Policy Enforcement Engine.

It is this specialized memory area that helps the ProVision ASIC achieves wire-speed performance when processing ACLs for packets. In fact, multiple passes through the TCAM can be performed for packet sizes that are typically found in customers' production networks. For the typical network, the average packet size will tend to be about 500 bytes. When maximum lookups are enabled, the ProVision ASIC performance is optimal for an average packet length of 200 bytes or more, which includes the range of packet sizes in typical networks.

The TCAM can support approximately 3,000 data entries that may be used to represent various traffic controls, including ACLs. For most customers, this quantity of entries will be more than adequate to ensure wire-speed performance for ACL processing. Keep in mind that each ACL entry may consist of multiple criteria such as a specific IP address and TCP or UDP port number.

In the initial release, the contents of the TCAM are common among the multiple line interface modules that a switch may have installed. For example, a ProCurve Switch 5406zl may have up to 6 line interface modules and a ProCurve Switch 5412zl may have up to 12 line interface modules.

Appendix C: Power over Ethernet

The IEEE 802.3af standard for Power over Ethernet (PoE) enables Ethernet switches to provide power, as well as network signals, for compliant devices over existing CAT-5 cable. Implementing devices that support integrated PoE, such as the ProCurve Switch 5400zl and 3500yl series, has several benefits:

- Lower cost installations because a single cable provides power along with data connectivity
- No need to modify existing network cabling
- Provides a truly international standard for power distribution
- Enables remote reset and power-off capability

PoE device types

There are two types of devices that are defined for PoE implementations. The first type, called the *powered device* or PD, receives power from the second type, called the *power sourcing equipment* or PSE. Powered devices include any Ethernet device capable of receiving power through a data port such as IP telephones, IP cameras, PDAs, and tablet PCs.

Power sourcing equipment, such as ProCurve switches with PoE support, must meet IEEE 802.3af specifications for voltage (47 to 57 volts DC) and wattage (up to 15.4 watts), with further limitations on the devices that receive power.

ProCurve switches that are PoE-enabled support both IEEE 802.3af-compliant devices as well as some pre-802.3af standard devices.

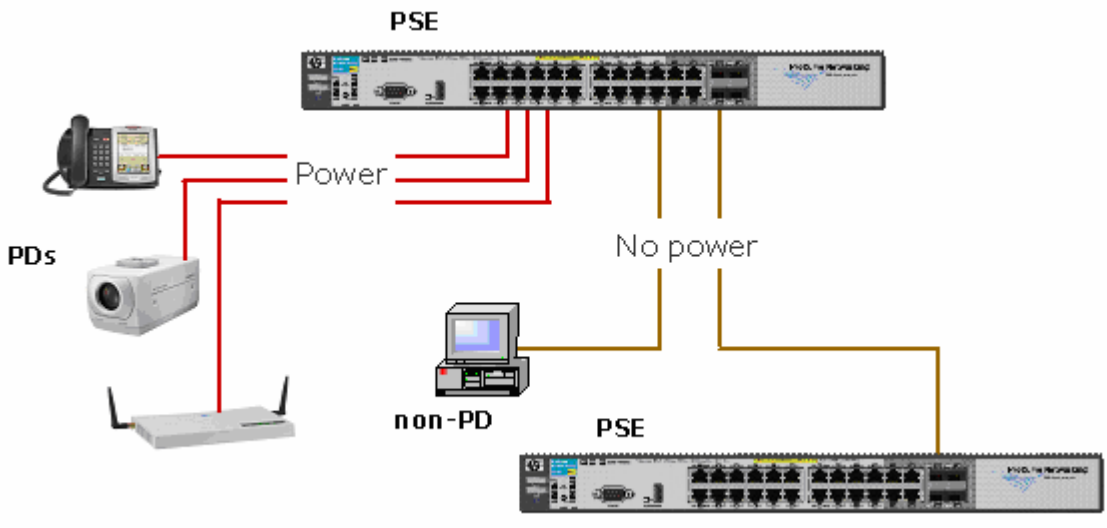

Power delivery options

The 802.3af standard provides two options for providing power over CAT 5 Ethernet cable. They are:

1. Unused pairs—This option takes advantage of the fact that 10Base-T and 100Base-TX signals use only two of the four twisted pairs in the cable. In this option, the pins on pairs 1/2 and 3/6 are used for data, and the pins on pairs 4/5 and 7/8 are used for power. The 4/5 pair is twisted together to form the positive supply, and the 7/8 pair is twisted together to form the negative supply.
2. Data pairs—This option provides power over pairs that also provide data and is necessary because 1000Base-T signals use all four pairs. In this option, either the 1/2 pair or the 3/6 pair can form the positive supply.

The 802.3af standard requires Power Source Equipment (PSE) to choose between the options. It is not possible to mix and match or to use all pairs for power transmission. ProCurve switches provide power over the data pairs. However, the standard requires PDs to be able to accept power over either data pairs or unused pairs.

PoE negotiation

In most networks, PSE will be connected to some devices that support PoE and some that do not. Consequently, in order to prevent damage to non-PoE devices, the 802.3af specification includes a negotiation mechanism between PSEs and the stations connected to them. Under the specification, PSEs apply a slight voltage on the power-delivery pairs. If there is a valid PD on the cable, it will present a resistance and a capacitance.

Typically, this discovery process requires approximately 500 milliseconds. The PSE will apply the full wattage if it discovers a PD, but the flow is current-limited to prevent damage to cable and equipment under fault conditions.

After discovery, a PD must draw a minimum current for the power flow to remain active. If, for example, the PD is unplugged, the PSE will discontinue power supply over the cable. The discovery process will be repeated when the PD is returned to service.

The ProCurve switches described in this document that are PoE enabled are:

- ProCurve Switch 5406zl—Supports up to 144 PoE ports.
- ProCurve Switch 5412zl—Supports up to 288 PoE ports.
- ProCurve Switch 3500yl-24G-PWR—Supports up to 24 PoE ports.

- ProCurve Switch 3500yl-48G-PWR—Supports up to 48 PoE ports.

Additional PoE Power – External Supplies

The ProCurve 620 Redundant/External Power Supply can be used with up to two Switch 3500s as an external supply providing an extra 398 W of PoE power for each switch. This extra power can be used to provide full PoE power redundancy at 15.4 W for all ports for the 3500yl-24G-PWR. It can also be used to provide full 15.4 W to all 48 ports on the 3500yl-48G-PWR, or full PoE power redundancy at 15.4 W for 24 ports on the 48 port switch.

The ProCurve Switch zl Power Supply Shelf provides a 3U high enclosure to hold up to two extra zl switch power supplies to provide additional PoE power to zl switches beyond the PoE power available through the switch's internal power supplies alone. These extra supplies can also be used to provide PoE power redundancy beyond what is possible with switch internal supplies. Large Switch zl VoIP installations are ideal environments for the ProCurve Switch zl Power Supply Shelf.

Support for Pre-802.3af Standard Powered Devices

The ProCurve Switch 5400zl and 3500yl series also support a number of pre-802.3af standard Powered Devices. These include:

- Cisco 7902G, 7905G, 7912G, 7940G, 7960G IP Phones
- Cisco Aironet 350, 1100, 1200, 1230AG Access Points

The Cisco pre-802.3af standard IP phones require the use of a straight-through cable for connectivity to the PSE. These IP phones are dependent on the polarity of the DC voltage (a violation of the IEEE 802.3af standard) and will not work properly if the Data Pairs (wires 1 and 2 and 3 and 6) are crossed over anywhere in the cable plant between the IP phone and Power Source Equipment.

Appendix D: PIM Sparse Mode

In Protocol Independent Multicast-Sparse Mode (PIM-SM), the assumption is that there are no hosts that want the multicast traffic unless they specifically ask for it. In contrast, PIM Dense Mode (PIM-DM) assumes downstream router membership unless it receives explicit prune message. PIM-SM is appropriate for wide scale deployment for both densely and sparsely populated groups and is the optimal choice for all production networks regardless of size and membership density.

The operation of PIM-SM centers on the use of a shared tree, with a router functioning as a Rendezvous Point (RP), the root of the tree. A shared tree prevents each router from maintaining source and group state information for every multicast source. Regardless of the number or location of multicast receivers, multicast senders register with the RP and send a single copy of multicast data through it to the registered receivers. Also, regardless of the location or number of sources, group members register to receive data and always receive it through the RP.

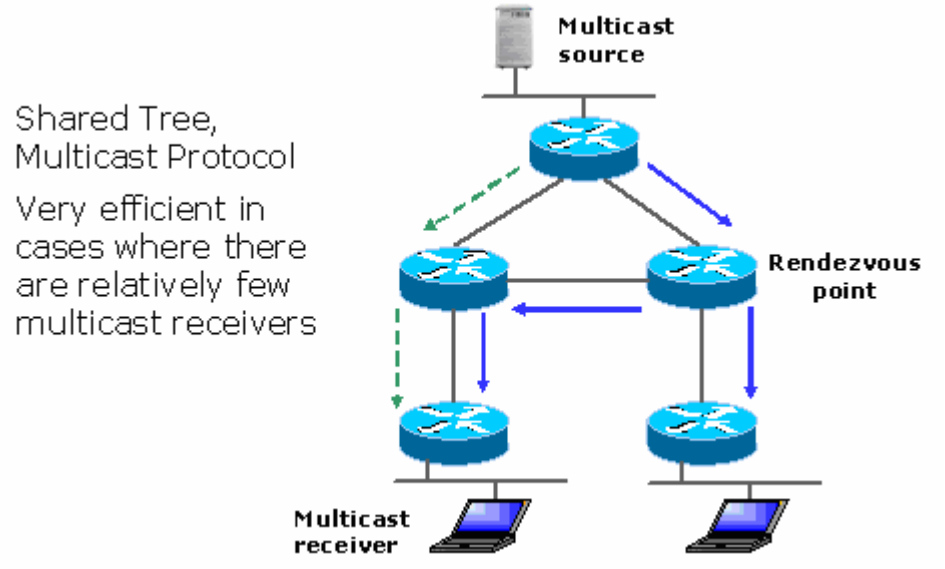

In order to receive a multicast stream, routers explicitly join the stream by sending join messages to the RP. This join message is analogous to a unicast router following a default route to a destination. Effectively, the function of the RP is a place for multicast sources and receivers to meet.

PIM-SM is extremely memory and CPU efficient. Since the only thing that most routers need to know is how to reach the RP, memory requirements are greatly reduced. There are several methods that can be used by routers in a PIM-SM domain to learn where to find the RP. Probably the simplest mechanism is statically configuring all routers to reach the RP. However, if the routers are statically configured to an RP and the RP fails, then the multicast network is no longer functional.

Alternatively, the RP can be learned dynamically through the PIM-SM bootstrap mechanism. Since this bootstrap mechanism is dynamic, it allows for network changes and redundancy. The PIM-SM bootstrap mechanism is generally the recommended approach for simplicity and redundancy.

Appendix E: LLDP-MED

IEEE 802.1AB Link Layer Discovery Protocol (LLDP) provides a standards-based method that enables devices such as ProCurve switches to advertise themselves to adjacent devices and to learn about adjacent LLDP devices. An LLDP packet transmitted by a ProCurve switch contains data about the switch and port. The switch advertises itself to adjacent (neighbor) devices by transmitting LLDP data packets out all ports on which outbound LLDP is enabled, and reading LLDP advertisements from neighbor devices on ports that are inbound LLDP-enabled. An LLDP-enabled port receiving LLDP packets inbound from neighbor devices stores the packet data in a Neighbor database (MIB).

LLDP provides the base capabilities for network devices, but was not considered sufficient for IP telephony devices. As a result, in 2004, an initiative by Mitel, ProCurve, Avaya and Enterasys was undertaken to enhance LLDP so that it could better support IP telephony devices. The development of LLDP-Medium Endpoint Discovery (LLDP-MED) (ANSI/TIA-1057/D6) extended the LLDP standard to support advanced features on the network edge for VoIP endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. The extensions to LLDP include the specification of additional TLV (Type, Length, and Value) entries specifically for VoIP management.

LLDP-MED benefits include:

- Plug-and-play provisioning for MED-capable, VoIP endpoint devices.
- Simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network.

- Automatic deployment of convergence network policies that includes voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority.
- Configurable endpoint location data to support the Emergency Call Service (ECS) such as Enhanced 911, 999 and 112.
- Detailed VoIP endpoint data inventory readable via SNMP from the switch.
- Power over Ethernet (PoE) status and troubleshooting support via SNMP.
- Support for IP telephony network troubleshooting of call quality issues via SNMP.

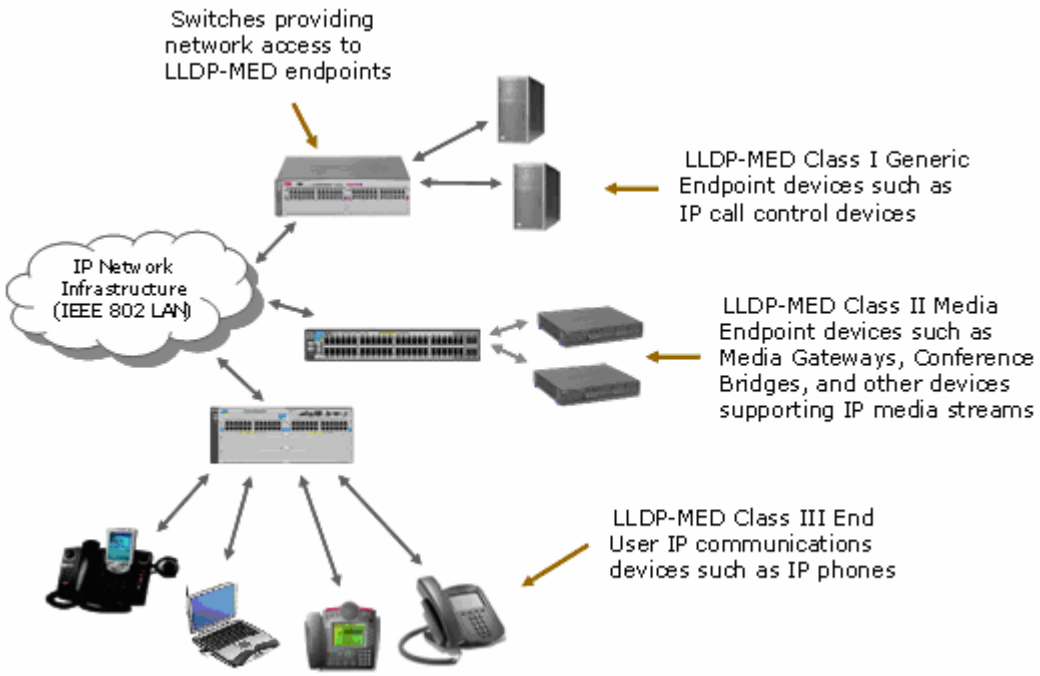
Switches providing network access to LLDP-MED endpoints
LLDP-MED Class I Generic Endpoint devices such as IP call control devices
IP Network Infrastructure (IEEE 802 LAN)
LLDP-MED Class II Media Endpoint devices such as Media Gateways, Conference Bridges, and other devices supporting IP media streams
LLDP-MED Class III End User IP communications devices such as IP phones

LLDP-MED endpoint devices are located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- Class 1 (Generic Endpoint Devices): These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- Class 2 (Media Endpoint Devices): These devices offer all Class 1 features plus media streaming capability, and include such devices as voice/media gateways, conference bridges, and media servers.
- Class 3 (Communication Devices): These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

Appendix F: Virus Throttle Security

Virus-throttle is based on the detection of anomalous behavior of network traffic that differs from a normal activity. Under normal activity, a computer will make fairly few outgoing connections to new computers, but instead is more likely to regularly connect to the same set of computers. This is in contrast to the fundamental behavior of a rapidly spreading worm, which will attempt many outgoing connections to new computers. For example, while computers normally make approximately one connection per second, the SQL Slammer virus tries to infect more than 800 computers per second.

Virus-throttle works by intercepting IP-routed connection requests, that is, connections crossing VLAN boundaries, in which the source subnet and destination subnet are different. The virus

throttle tracks the number of recently made connections. If a new, intercepted request is to a destination to which a connection was recently made, the request is processed as normal. If the request is to a destination that has not had a recent connection, the request is processed only if the number of recent connections is below a pre-set threshold. The threshold specifies how many connections are to be allowed over a set amount of time, thereby enforcing a connection rate limit. If the threshold is exceeded, because requests are coming in at an unusually high rate, it is taken as evidence of a virus. This causes the throttle to stop processing requests and, instead, to notify the system administrator.

This applies to most common Layer 4 through 7 session and application protocols, including TCP connections, UDP packets, SMTP, IMAP, Web Proxy, HTTP, SSL, and DNS—virtually any protocol where the normal traffic does not look like a virus spreading. For virus-throttle to work, IP routing and multiple VLANs with member ports must first be configured.

Note that some protocols, such as NetBIOS and WINS, and some applications such as network management scanners, notification services and p2p file sharing are not appropriate for virus-throttle. These protocols and applications initiate a broad burst of network traffic that could be misinterpreted by the virus-throttle technology as a threat.

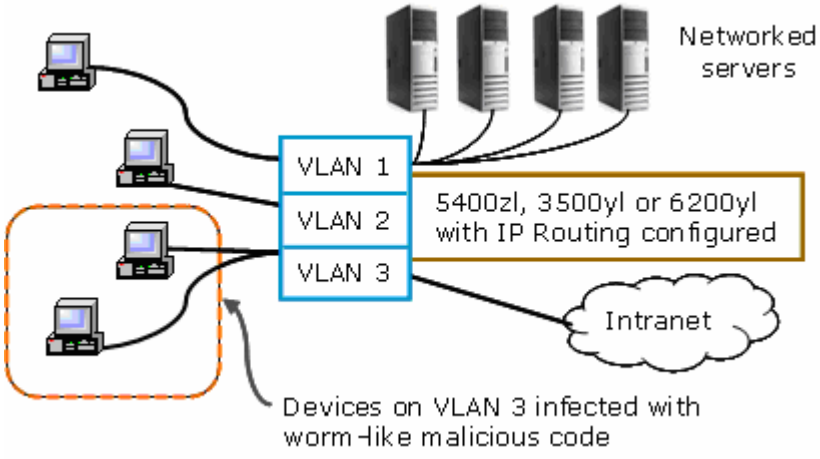

On the ProCurve Switch 5400zl, 3500yl, and 6200yl series, virus throttle is implemented through connection-rate filtering. When connection-rate filtering is enabled on a port, the inbound routed traffic is monitored for a high rate of connection requests from any given host on the port. If a host appears to exhibit the worm-like behavior of attempting to establish a large number of outbound IP connections in a short period of time, the switch responds depending on how connection-rate filtering is configured.

Response options

The response behavior of connection-rate filtering can be adjusted by using Filtering options. When a worm-like behavior is detected, the connection-rate filter can respond to the threats on the port in the following ways:

- **Notify only of potential attack:** While the apparent attack continues, the switch generates an Event Log notice identifying the offending host source address (SA) and (if a trap receiver is configured on the switch) a similar SNMP trap notice.
- **Notify and reduce spreading:** In this case, the switch temporarily blocks inbound routed traffic from the offending host source address for a “penalty” period and generates an Event Log notice of this action and a similar SNMP trap notice if a trap receiver is configured on the switch. When the penalty period expires the switch re-evaluates the routed traffic from the host and continues to block this traffic if the apparent attack continues. During the re-evaluation period, routed traffic from the host is allowed.
- **Block spreading:** This option blocks routing of the host‘s traffic on the switch. When a block occurs, the switch generates an Event Log notice and a similar SNMP trap notice if a trap receiver is configured on the switch. Note that system personnel must explicitly re-enable a host that has been previously blocked.

Sensitivity

The ability of connection-rate filtering to detect relatively high instances of connection-rate attempts from a given source can be adjusted by changing the global sensitivity settings. The sensitivity can be set to low, medium, high or aggressive as described below:

- **Low:** Sets the connection-rate sensitivity to the lowest possible sensitivity, which allows a mean of 54 routed destinations in less than 0.1 seconds, and a corresponding penalty time for Throttle mode (if configured) of less than 30 seconds.
- **Medium:** Sets the connection-rate sensitivity to allow a mean of 37 routed destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 30 and 60 seconds.
- **High:** Sets the connection-rate sensitivity to allow a mean of 22 routed destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 60 and 90 seconds.
- **Aggressive:** Sets the connection-rate sensitivity to the highest possible level, which allows a mean of 15 routed destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 90 and 120 seconds.

Connection-rate ACL

Connection-rate ACLs are used to exclude legitimate high-rate inbound traffic from the connection-rate filtering policy. A connection-rate ACL, consisting of a series of access control entries, creates exceptions to these per-port policies by creating special rules for individual hosts, groups of hosts, or entire subnets. Thus, the system administrator can adjust a connection-rate filtering policy to create and apply an exception to configured filters on the ports in a VLAN.

Appendix G: VRRP

Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in the static default routed environment. In a VRRP environment, two or more “virtual” routers cooperate to provide a high availability capability on a LAN. VRRP specifies an election protocol that dynamically assigns routing responsibility to one of the virtual routers on a LAN.

A virtual router consists of a set of router interfaces on the same network that share a virtual router identifier (VRID) and a virtual IP address. One router in the group becomes the VRRP Master and the other routers are designated as VRRP Backups. The VRRP Master controls the IP addresses associated with a virtual router.

The VRRP Master router periodically sends advertisements to a reserved multicast group address. The VRRP Backup routers listen for advertisements and one of the backups will assume the Master role, if necessary. A VRRP router can support many virtual router instances, each with a unique VRID/IP address combination. The election process provides dynamic fail-over to one off the remaining VRRP Backups should the Master become unavailable.

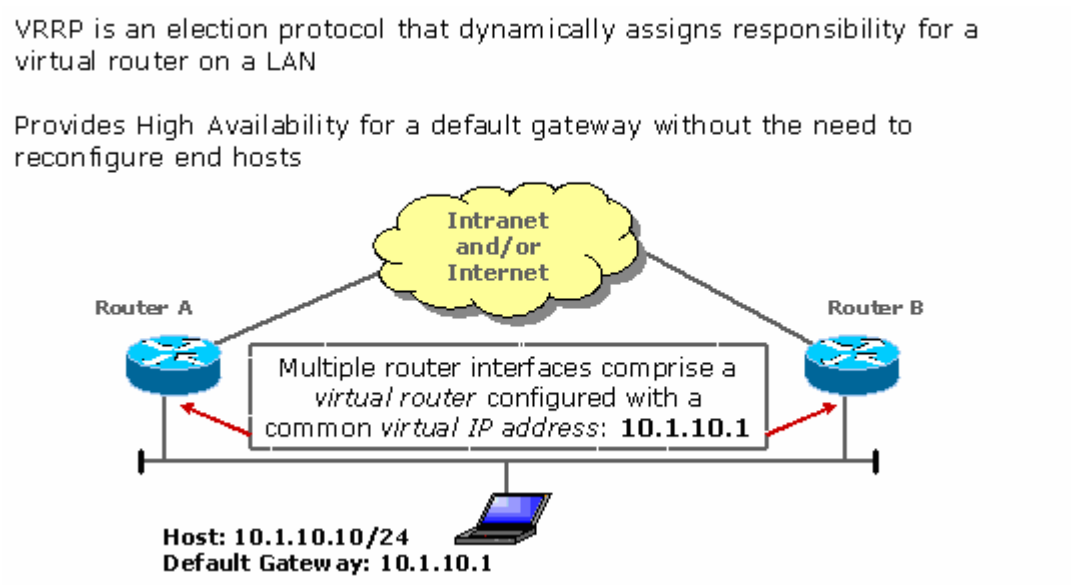
VRRP is an election protocol that dynamically assigns responsibility for a virtual router on a LAN
Provides High Availability for a default gateway without the need to reconfigure end hosts
Intranet and/or Internet
Router A
Router B
Multiple router interfaces comprise a virtual router configured with a common virtual IP address: 10.1.10.1
Host: 10.1.10.10/24
Default Gateway: 10.1.10.1

The virtual IP address shared by a group of VRRP routers on a given network segment functions as the next hop IP address used by neighboring hosts. The VRRP Master router simply forwards packets that have been received from hosts using the VRRP Master as the next hop gateway. The existence of a VRRP master and one or VRRP Backups is transparent to the neighboring hosts.

The advantage gained from using VRRP is a default path with higher availability, but without requiring configuration of dynamic routing or router discovery protocols on every end-host.

VRRP on ProCurve switches is interoperable with other routers that support RFC 3768.

VRRP operational aspects include:

- Pre-emptive mode can be disabled to prevent VRRP router flapping
- Default Advertisement interval of 1 sec
- Default Detection time of 3.6 seconds

XRRP support on 5300xl switch

Similar in concept to VRRP, the ProCurve XL Router Redundancy Protocol (XRRP), provides the equivalent mechanism in the ProCurve Switch 5300xl series for backup functionality. Like VRRP, XRRP presents a virtual router to the end node connections whose IP and MAC address can transition from the master ProCurve Switch 5300xl to the backup ProCurve Switch 5300xl in the event the master ProCurve Switch 5300xl interface fails.

XRRP does not interoperate with VRRP, but can coexist in a VRRP environment without interference.

Appendix H: OSPF Equal Cost Multipath

In Open Shortest Path First (OSPF), if different subnet destinations in a network are reachable through multiple equal-cost, next-hop routes, the router chooses the same next-hop route at a given point in time to send traffic to destinations reachable through that next-hop router.

With OSPF Equal Cost Multipath (OSPF-ECMP), routers support optional load-sharing across redundant links where the network offers two or more equal-cost next-hop routes for traffic to different subnets. All traffic for different hosts in the same subnet goes through the same next-hop router. Multiple paths are balanced based on the number of destination subnets.

ProCurve's OSPF-ECMP feature is interoperable with OSPF-ECMP implementations from various vendors including Cisco, 3Com and Extreme switches. The ProCurve implementation supports up to four ECMP links.

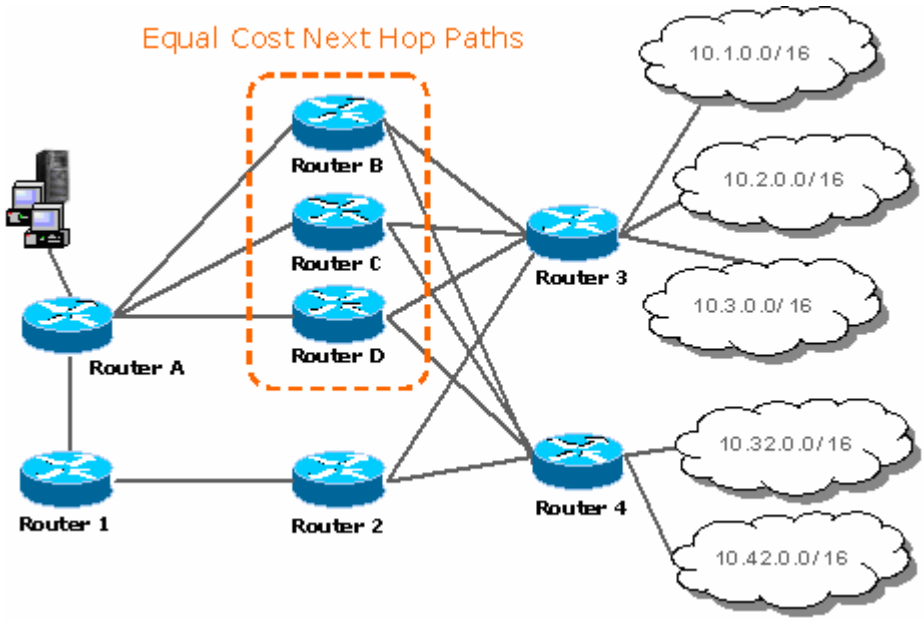

The example shows that there are three equal-cost, next-hop paths from Router A to the destination subnets with load-sharing across redundant links. At any point in time, Router A's routing table could have information indicating the following:

Destination Subnet	Next Hop
10.1.0.0/16	Router C
10.2.0.0/16	Router D
10.3.0.0/16	Router B
10.32.0.0/16	Router B
10.42.0.0/16	Router D

Appendix I: Troubleshooting

LED Status Indicators for 5400zl

The ProCurve Switch 5400zl series management module has various LED status indicators that are described in the following table.

LED	State	Indication
Power	On (green)	The switch is receiving power.
	Off	The switch is NOT receiving power.
Fault	On (orange)	On briefly at the beginning of switch self test after the switch is powered on or reset. If on for a prolonged time, the switch has encountered a fatal hardware failure, or has failed its self test.
	Off	The normal state; indicates there are no fault conditions on the switch.
	Blinking (orange)	A fault has occurred on the switch, one of the switch modules, an individual port, a power supply, or a fan. The Status LED for the module or other device with the fault will flash simultaneously.

LED	State	Indication
Test	On (green)	The switch self test and initialization are in progress after you have power cycled or reset the switch. The switch is not operational until this LED goes off. The Self Test LED also comes on briefly when you “hot swap” a module into the switch and the module is automatically self tested.
	Off	Normal operation; the switch is not undergoing self test.
	Blinking (orange)	A component of the switch has failed its self test. The status LED for that component, for example a switch module, and the switch Fault LED will flash simultaneously.
DIMM	On (green)	Normal operation.
	Off	DIMM status is unknown.
	Blinking (orange)	If DIMM, Fault, and Self Test LEDs are blinking, DIMM failed self-test. If DIMM and Fault LEDs are blinking, an operational fault has occurred. If blinking at a fast rate, an operational alert occurred and is unresolved.
Flash	On (green)	Flash Card status is known and fault free
	Off	Flash Card status is unknown.
	Blinking (orange)	If Flash, Fault, and Self Test LEDs are blinking, Flash Card failed self-test. If Flash and Fault LEDs are blinking, an operational fault has occurred. If blinking at a fast rate, an operational alert occurred and is unresolved. For example, Flash Card is not present.
Mgmt	On (green)	A Management module is present and fault free.
	Off	The switch is powered off.
	Blinking (orange)	There is a fault on the Management module.

LED	State	Indication
PoE	On (green)	If any PoE modules are installed.
	Off	If no PoE modules are installed.
	Blinking (orange)	Internal PoE fault, external load fault, or denied PoE power.
Temp	On	Internal temperature is normal.
	Blinking (orange)	An over temperature condition has been detected.
Fan	On (green)	Normal operation.
	Blinking (orange)	One or more of the switch’s fans have failed. The switch Fault LED will be blinking simultaneously.
Internal Power (numbers corresponding to the power supply positions)	On (green)	A power supply is installed in the position in the back of the switch corresponding to the number, and the supply is plugged in to an active AC power source. As shipped, the switch has a single power supply in position 1.
	Off	A power supply is not installed in the position corresponding to the number.
	Blinking (orange)	The power supply installed in the position corresponding to the number is not plugged in to an active AC power source, or has experienced a fault. The switch Fault LED will be blinking simultaneously.

EPS	On (green)	An external power supply is connected.
	Off	An external power supply is not connected.
	Blinking (orange)	The external power supply has a fault, or is connected but not plugged into AC power.
Modules A, B, C, D, E, F, G, H, I, J, K, L	On (green)	A module is installed in the switch module slot corresponding to the letter and the module is undergoing or has passed self test. This also occurs when you install a module when the switch is already powered on (“hot swap”).
	Off	A module is not installed in the switch module slot corresponding to the letter.
	Blinking (orange)	The module status LED flashes very briefly when a module is being hot swapped. If the LED flashes for a prolonged time, the module in the slot corresponding to the letter has failed self test or encountered some other fault condition.
Modules in PoE Mode	On (green)	PoE is normal for this slot.
	Off	The module in this slot is not a PoE module.
	Blinking (orange)	If any port has an internal hardware failure, blinks once every 1.6 seconds. If any port is denied PoE power or detecting an external PD fault, blinks once every 0.5 seconds.

LED	State	Indication
LED Mode Select	Act (green)	Flickers to show relative activity.
	FDx (green)	Flickers to show relative activity.
	PoE (green)	Indicates which ports are supplying power. • If the Mode LED is on the port is providing PoE power. • If the Mode LED is off the port is not providing PoE power. • If the Link LED is on the port is enabled for PoE. • If the Link LED is off the port is disabled for PoE. • If the Link LED is blinking, the port has an error or the port is denied power due to insufficient power.
	Spd (green)	Indicates speed of operation of each port. • If the Port LED is off, the port is operating at 10 Mbps • If the Port LED is blinking, the port is operating at 100 Mbps • If the Port LED is on continuously, the port is operating at 1000 Mbps
	Usr (green)	Reserved for future development.

Table 10. LED Status Indicators for ProCurve Switch 5400zl Series—Management Module

The LED status indicators for the ProCurve Switch 5400zl series switch modules are described in the following table.

LED	State	Indication
Link	On (green)	Indicates the port is enabled and receiving a link beat signal (for the twisted-pair ports), or a strong enough light level (for the fiber-optic ports) from the connected device.
	Off	Indicates no active network cable is connected to the port, the port is not receiving link beat or sufficient light, or the port has been disabled.
	Blinking (orange)	If the port LED is blinking simultaneously with the Fault LED, the corresponding port has failed its self test.

Mode	Depending on the mode selected, displays: • Network activity information • Whether the port is configured for Full Duplex operation • Maximum speed operation • Whether PoE power is being supplied or not

Table 11. LED Status Indicators for ProCurve Switch 5400zl Series—Switch Modules

LED Status Indicators for 3500yl and 6200yl

The front panel of the ProCurve Switch 3500yl and 6200yl series has various LED status indicators that are described in the following table.

LED	State	Indication
Power	On (green)	The switch is receiving power.
	Off	The switch is NOT receiving power.
Fault	On (orange)	On briefly after the switch is powered on or reset, at the beginning of switch self test. If this LED is on for a prolonged time, the switch has encountered a fatal hardware failure, or has failed its self test.
	Off	The normal state; indicates there are no fault conditions on the switch.
	Blinking (orange)	A fault has occurred on the switch, one of the switch ports, module in the rear of the switch, or the fan. The status LED for that component will blink simultaneously.
Test	On (green)	The switch self test and initialization are in progress after the switch has been power cycled or reset. The switch is not operational until this LED goes off. The Self Test LED also comes on briefly when you “hot swap” a mini-GBIC into the switch; the mini-GBIC is self tested when it is hot swapped.
	Off	The normal operational state; the switch is not undergoing self test.
	Blinking (orange)	A component of the switch has failed its self test. The status LED for that component will blink simultaneously.
Temp	On (green)	Internal temperature is normal.
	Blinking (orange)	An over temperature condition has been detected.
Fan Status	On (green)	Normal operation.
	Blinking (orange)	One of the unit’s fans has failed. The switch Fault LED will be blinking simultaneously.
PoE Status	On (green)	If any ports are supplying PoE power.
	Off	If no ports are supplying PoE power. Should be off only during the boot process.
	Blinking (orange)	If any port has an internal hardware failure, blinks once every 1.6 seconds. If any port is denied PoE power or detecting an external PD fault, blinks once every 0.5 seconds.

LED	State	Indication
RPS Status	On (green)	Normal operation. RPS is connected and operating correctly. RPS could be powering the unit.
	Off	RPS is not connected or is not powered on.

LED	State	Indication
	Blinking (orange)	RPS is connected, but has experienced a fault.
EPS Status	On (green)	Connected to an EPS, and receiving power.
	Off	The switch is not connected to an EPS.
	Blinking (orange)	The EPS is connected but has experienced a fault or is unplugged.
Mdl—Module Status	On (green)	Expansion module is plugged into expansion slot and operating correctly.
	Off	Expansion module is not plugged into expansion slot.
	Blinking (orange)	Expansion module is plugged into expansion slot but has experienced a fault.
Port LEDs (Link and Mode)	Link (green)	Indicates the port LEDs are displaying link information. • If the port LED is on, the port is enabled and receiving a link indication from the connected device. • If the port LED is off, the port has no active network cable connected, or is not receiving link beat or sufficient light. Otherwise, the port may have been disabled through the switch console, the web browser interface, or ProCurve Manager. If the port LED is blinking (orange) simultaneously with the Fault LED, the corresponding port has failed its self test.
	Mode (green)	The operation of the Mode LED is controlled by the LED Mode select button, and the current setting is indicated by the LED Mode indicator LEDs near the button. Press the button to step from one view mode to the next. The default view is Activity (Act).
LED Mode Select	Act (green)	Flickers to show relative activity.
	FDx (green)	Flickers to show relative activity.
	PoE (green)	Indicates which ports are supplying power. • If the Mode LED is on the port is providing PoE power. • If the Mode LED is off the port is not providing PoE power. • If the Link LED is on the port is enabled for PoE. • If the Link LED is off the port is disabled for PoE. • If the Link LED is blinking, the port has an error or the port is denied power due to insufficient power.
	Spd (green)	Indicates speed of operation of each port. • If the Port LED is off, the port is operating at 10 Mbps • If the Port LED is blinking, the port is operating at 100 Mbps • If the Port LED is on continuously, the port is operating at 1000 Mbps
	Usr (green)	Reserved for future development.

Table 12. LED Status Indicators for ProCurve Switch 3500yl and 6200yl Series

Part Numbers and Field Replaceable Units

The table below lists various part numbers and Field Replaceable Units (FRUs) for the ProCurve Switch 5400zl series.

Part Number	Component
J8702-69001	ProCurve Switch 5400zl 24 port Gig-T PoE module

J8705-69001	ProCurve Switch 5400zl 20 port Gig-T + 4 port mini-GBIC module
J8707-69001	ProCurve Switch 5400zl 4 port 10G X2 module
J8726-69001	Management Module without CF card
J8697-60005	Fan Tray Assembly for 5406zl
J8698-60005	Fan Tray Assembly for 5412zl
5070-1056	Programmed CF kit for Management Module
J8713-69001	High-Power PoE Supply 1500W
J8712-69001	875W Power Supply
5184-1894	Console cable
5069-8561	Two-post Rack Mounting Kit for 5406zl
5069-8562	Two-post Rack Mounting Kit for 5412zl
356578-B21	Four-post Rack Rail Mounting Kit for 5406zl and 5412zl

Table 13. Part Numbers and FRUs for ProCurve Switch 5400zl Series

The switches normally ship with a rack mounting kit that allows installation into a two-post, 19-inch data communications rack. If installation into a 10K four-post 19-inch rack or a four-post rack meeting the standard EIA unit of measurement is desired, a rack mounting kit is available providing rails that give sturdy support for the switch along its entire length.

The table below lists various part numbers and FRUs for the ProCurve Switch 3500yl and 6200yl series.

Part Number	Component
J8692-69001	ProCurve Switch 3500yl-24G-PWR Intelligent Edge
J8693-69001	ProCurve Switch 3500yl-48G-PWR Intelligent Edge
J8992-69001	ProCurve Switch 6200yl-24G mini-GBIC
J8694-69001	ProCurve Switch 3500yl/6200yl 2p 10GbE X2 + 2p CX4 Module
J8696-69001	ProCurve 620 Redundant/External Power Supply
5069-8589	Fan Tray Assembly for 3500yl and 6200yl
5184-1894	Console cable
5069-5705	Two-post Rack Mounting Kit
356578-B21	Four-post Rack Rail Mounting Kit

Table 14. Part Numbers and FRUs for ProCurve Switch 3500yl and 6200yl Series

The table below lists the part numbers of power cords that can be used with the ProCurve Switch 5400zl, 3500yl and 6200yl series products. The power cords available for these switches are sized for the increased current that can be drawn meeting the needs of PoE. As a result, these power cords may not be found in a typical environment and power cords “borrowed” from other products will not work in most instances.

Country	ProCurve Switch 3500yl and 6200yl	ProCurve Switch 5400zl	
		875W Power Supply	1500W Power Supply
UK, Hong Kong, Singapore	8120-5334	8120-5334	8120-6898
Australia, New Zealand	8120-5335	8120-5335	8121-0871
Europe	8120-5336	8120-5336	8120-6899
US, Canada, Mexico	8121-0914	8121-0973	8121-6903
Japan	8120-5342	8120-5342	8120-6903
China	8120-1034	8120-8385	8120-0924
South Africa, India	8120-5341	8120-5341	8120-0915
Switzerland	8121-5339	8121-5339	8121-0915
Denmark	8120-5340	8120-5340	8120-6897

Table 15. Power Cords for ProCurve Switch 5400zl, 3500yl and 6200yl Series

To find out more about ProCurve Networking products and solutions, visit our web site at

www.hp.com/go/procurve

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA0-5388ENW, Revision 1 01/2007